Noise-Robust Modes of the Retinal Population Code Geometrically Correspond with "Ridges"

Adrianna R. Loback[1*], Jason S. Prentice[1], Mark L. Ioffe[2], Michael J. Berry II[1,3]

1 *Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA*
2 *Department of Physics, Princeton University, Princeton, NJ 08544, USA*
3 *Department of Molecular Biology, Princeton University, Princeton, NJ 08544, USA*
*adrianna@princeton.edu

(Dated: October 24, 2016)

arXiv:1610.06886v1 [q-bio.NC] 21 Oct 2016

Throughout the brain, local circuits encode information using large populations of neurons, yet the structure of this population code remains poorly understood. Previous studies applying the maximum entropy principle to retinal multi-electrode array data have shown that the joint probability distribution over all retinal ganglion cell population states is proliferated by local maxima, which represents a form of clustering of visual stimuli into a discrete set of categories. However, these studies focused on responses to highly repeated stimuli. Another recent study using a hidden Markov model approach likewise found that ganglion cell population activity partitioned into noise-robust clusters, and further that this clustering was present for the non-repeated stimulus regime. Here, we investigate the relationship between the clusters found using these two different modeling approaches. We first demonstrate that local maxima in the probability landscape result from the repeat structure of the stimulus ensemble, and are absent in the more realistic, non-repeated stimulus regime. Nonetheless, we show that structure is still present in the non-repeated stimulus regime in the form of what we call "soft local maxima," which are local maxima within the subspace of all neural activity states that are restricted to have the same spike count. Using a numerical approach, we show that soft local maxima are organized into a reliable sequence across the different spike count levels, which altogether has the geometry of a "ridge." These ridges closely corresponded with the clusters formulated by the hidden Markov model, and moreover with neuron "communities," in the sense defined in computer science. These findings suggest a unified description of the ganglion cell population code across different stimulus repeat statistics, and suggest that ridges or communities may be an important feature of the output population activity extracted by visual processing areas downstream of the retina.

While the manner in which information is represented in the brain has been intensively studied at the single neuron level [1–3], ample experimental evidence and theoretical consideration have made it evident that neural circuits generally encode and transmit information using large populations of neurons [4–17]. With the advent of new experimental technologies, it is becoming increasingly tractable to record from large neuron populations [18–21], facilitating questions about the empirical nature of neural population coding that were previously unaddressable. Fundamental to most conceptual approaches toward answering these questions is, as with the single neuron case, characterizing the probability distribution over all neural responses. The key additional issue when considering the multi-neuron scenario however, is the nature of correlations among neurons, which fundamentally shape the joint probability distribution of population response patterns. Recent information theoretic and statistical model-based approaches applied to data have established that pairwise and higher-order correlations are highly non-negligible in real, large neuron populations [22–26, 28, 38]. Moreover, analytical studies have shown that the structure of correlations has a profound effect on shaping the capacity of neural population codes [29–32].

Due to the combinatorial explosion of possible joint responses - i.e. a population of N neurons, with binarized spiking, could theoretically encode 2^N distinct binary population response patterns - and limited sampling, it is generally empirically intractable to determine the joint probability mass function (p.m.f.) for $N > 20$ neurons. This is problematic, given that the relevant population coding units in many brain regions are of the order of hundreds of neurons. For example, in the salamander retina, it is known that a single region of visual space maps onto a population of retinal ganglion cells (RGCs) numbering ∼200 cells [20]. To give another example, an average single pyramidal neuron receives synaptic inputs from tens of thousands of excitatory synapses, and several thousand inhibitory synapses [8]. Thus, in order to investigate questions about population coding at the relevant scale, a model of the joint p.m.f is needed.

Within computational neuroscience, there are two fundamentally distinct methodologies that have been used to create probabilistic models of network activity. The traditional approach - which includes linear-nonlinear-Poisson (LNP) and generalized linear models (GLMs) - attempts to explicitly capture the dependence of neuron responses on the stimulus [23, 33, 34]. Here we will re-

fer to such models as "encoding models," as this term has been previously used [34, 35]. The second approach, which we call the "activity model" approach, involves directly modeling the structure of population activity, without any reference to the stimulus. An important advantage of activity models is that they correspond with the unsupervised problem actually faced by downstream brain areas, which lack direct access to both the external stimuli and biophysical parameters of the upstream area. The most well-studied activity models are the class of Maximum Entropy (MaxEnt) models [22, 26, 36–42]. Conceptually, the central idea of the MaxEnt framework is to choose the unique model of the p.m.f. which reproduces a given set of experimental observations, while importantly being the least structured among all possible models that likewise satisfy these constraints. The well-studied Pairwise MaxEnt model takes as constraints the experimentally-estimated mean firing rates and pairwise correlations [22].

A notable feature of the Pairwise MaxEnt approach is that the mathematical form of the modeled probability distribution is equivalent to a well-known distribution in statistical physics: the Boltzmann distribution of an Ising spin glass model at a temperature $T = 1$ [43]. The pairwise interaction parameters (see Eq. (18) in *Methods*) learned by fitting MaxEnt models have been reported to feature both positive and negative values [41]. In spin glass models, such disordered interactions are known to lead to an energy function with many local minima, also called "basins" (see Fig 1B). Consistent with this, previous work which fit the K-Pairwise MaxEnt model (see *Methods*) to a large dataset of RGCs responding to a highly-repeated natural movie stimulus reported a proliferation of basins in the energy landscape [41]. Thus, arising from the Maximum Entropy literature was the idea that a proliferation of local maxima is a feature of the RGC population response probability landscape.

The reason why this feature is important, is due to its connections with another critical issue that must be considered when attempting to discern the nature of neural population codes: noise. Noise is prevalent at all levels throughout the nervous system [2, 44–47]. Consequently, even when presented with identical stimuli, sensory neuron populations exhibit joint response patterns with high trial-to-trial variability [17, 35, 41, 48–50]. How the brain solves the problem of accurately decoding population activity in the presence of noise is unknown. In communications engineering, a similar problem of achieving the high degree of reliability required by modern data transmission and storage systems has been addressed by the design of what are called "error-correcting codes", which are encoding schemes that enable accurate decoding even when the transmitted information has been corrupted by noise [51, 52]. Traditionally, engineered error-correcting codes achieve this robustness by introducing redundancy so that the full space of possible output patterns is partitioned into subsets, such that all noise-corrupted versions of the same input message map to the same subset. Consistent with this notion, it has been empirically found that the retinal code is highly redundant [12]. Likewise consistent, the authors in [41] found that the set of disparate population responses elicited by repeated presentations of an identical stimulus could be mapped to a basin assignment which was reproducible across the repeats. This noise-robust partitioning idea has also gained attention by other recent works in the field [17, 35, 53, 54].

In particular, using another type of activity model called the Tree Hidden Markov model (HMM), which was recently developed in [35], it was found that the joint responses of large populations of RGCs naturally partitions into a discrete set of "collective modes." Moreover, the many disparate response patterns evoked by

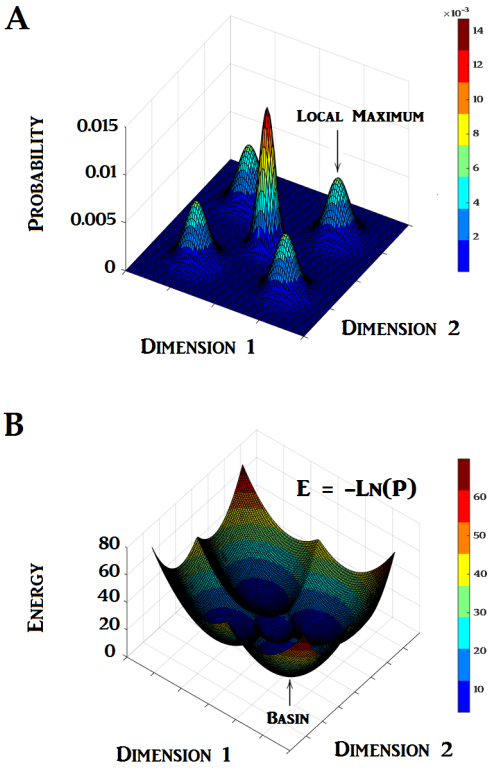

FIG. S1: **Schematic illustrating the concept of local maxima and basins. (A)** Cartoon illustrating local maxima in a probability landscape. Note that for ease of visualization, we depict the full space of population activity patterns as projected into a continuous, 2-dimensional space. The actual domain of the joint p.m.f. considered in the present paper is discrete and high-dimensional; in particular, an N-dimensional hypercube. **(B)** Cartoon of basins in the energy landscape corresponding with the probability landscape shown in panel (A). A response state $\vec{\sigma}$ which is a local maximum of the joint p.m.f. P equivalently corresponds with a basin in the associated energy landscape, since the energy $E(\vec{\sigma}) \equiv -\ln P(\vec{\sigma})$.

an identical stimulus were found to preferentially map to the same collective mode [35]. Importantly, these latent modes were identified by fitting the Tree HMM only to the portion of responses elicited by *non-repeated* stimulus segments. Thus, in contrast to [41], this error-correction capacity was obtained for the non-repeated stimulus regime. It should be emphasized that the collective modes in [35] were explicitly formulated by the Tree HMM statistical model, whereas the local maxima reported in [41] were defined geometrically after fitting the MaxEnt model. This thus leaves open the question: How are these two different definitions of population "codewords" related? More specifically, what, if any, is the precise geometric interpretation of the collective modes in relation to the RGC joint response probability landscape?

Here, we address this question. We obtained multi-electrode array data from large datasets of salamander RGCs simultaneously responding to various visual stimuli (population sizes ranged from 128 to 170 RGCs), and modeled the corresponding response probability landscapes with either the Tree HMM or K-Pairwise MaxEnt model. We found that the qualitative structure of the RGC joint response probability landscape depended strongly on the repeat structure of the stimulus ensemble. In particular, for non- and mildly-repeated stimulus ensembles, non-silent local maxima were exceedingly rare. We then introduced a new notion of structure, here termed "soft local maxima," which conceptually are local maxima in the space of all neural responses that are restricted to have the same spike count, K. These soft local maxima were robustly prevalent at each spike count level examined. Using a numerical approach, we then explored their organization across spike count levels. This approach leads to an algorithm that links together different soft local maxima into a discrete "ridge", which is a geometric feature of the joint probability landscape of neural activity. We found that there was a close correspondence between these ridges and the collective modes formulated by the Tree HMM. Our results moreover indicate that these ridges may correspond with neuron "communities," a notion from computer science [55]. Thus, our results suggest a unified picture linking statistically-defined noise-robust modes of the retinal population code and geometric structure in the joint probability landscape. Specifically, we find that under broad stimulus ensembles, neural population responses cluster into a discrete set of ridges; but when particular stimuli are encountered frequently, local peaks are additionally formed in the probability landscape.

RESULTS

Qualitative structure of the RGC population response probability landscape depends on the repeat structure of the stimulus

Given that previous studies that used the Maximum Entropy principle to explore the probability landscape of neural population activity have primarily used highly repeated stimulus ensembles, it was not previously addressed whether or how the result of a proliferation of local maxima in the modeled response probability landscape depends on the stimulus. In particular, how does this feature depend on the repeat structure of the stimulus ensemble? For a highly-repeated stimulus ensemble - such as the natural movie used in [41], which was comprised of 297 repeats of a 19 s movie clip - the exact same stimulus is sampled on many trials. It is thus plausible that the previously reported non-silent probability local maxima could correspond with the 'average' response pattern elicited by the same stimulus, with the width of the local maximum being attributable to noise. To investigate whether a proliferation of local maxima is a general feature which is independent of the stimulus repeat structure, we first applied the same K-Pairwise MaxEnt model used in [41] to a large, experimental dataset of the responses of $N = 128$ retinal ganglion cells to a mildly-repeated natural movie stimulus (see *Methods: Visual stimulus display*). We will refer to this stimulus ensemble throughout as Movie #1. Consistent with past work, we defined the joint activity pattern $\vec{\sigma}(t)$ as the binary N-dimensional vector of instantaneous states of each neuron in the population at time bin t: the ith entry σ_i being 1 if the ith neuron spiked, and 0 if it was silent. After fitting the model to the data (see *Methods: Fitting the K-Pairwise Maximum Entropy model*), we then searched for local maxima of the modeled probability landscape using the same numerical method as in [41], which we term "single spin flip ascent."

In practice, this involves taking each population activity pattern observed in the data, and moving "uphill" on the modeled joint probability landscape (see *Methods: Exploring the modeled probability landscape*). The termination patterns of this procedure are by definition local maxima of the modeled probability landscape, when only allowing the state of a single neuron to change (i.e. using a Hamming distance of one as the distance metric). This single spin flip procedure can be thought of as a mapping from the full joint response space of 2^N states to the space of local probability maxima.

As shown in Fig 2A, performing the single spin flip ascent procedure (using the probability landscape of joint activity obtained by fitting the K-Pairwise MaxEnt model) on the mildly-repeated stimulus dataset resulted in 65 unique local maxima. Of the 175,002 observed population responses in the data, 99.71% were mapped via

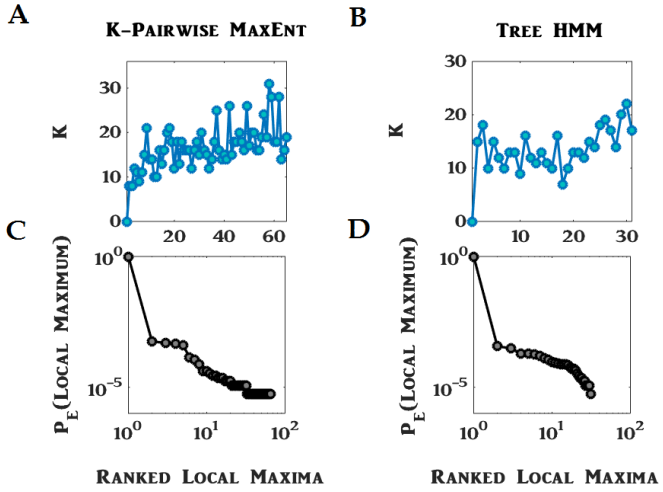

FIG. S2: **Local maxima results for a dataset of** 128 **retinal ganglion cells responding to Movie** #1**, a mildly-repeated stimulus. (A,C)** Results using the K-Pairwise Maximum Entropy model (8385 parameters). **(A)** The spike count, denoted by $K \equiv \sum_{i=1}^{N} \sigma_i$, for each of the 65 unique local maxima identified (x-axis, ranked). **(C)** Log-log plot of the proportion (y-axis, denoted by P_E) of the 175,002 population responses observed in the data which were mapped via single spin flip ascent to the corresponding unique local maximum indicated on the x-axis (ranked). **(B,D)** Results using the Tree Hidden Markov model (8379 parameters). The panels are as described in (A) and (C), but for the Tree HMM results.

single spin flip ascent to the all-silent local maximum (i.e. $\vec{0}$), and 0.037% were mapped to the rank-2 local maximum. This is in stark contrast to the $\sim 50\%$ of responses which were mapped via single spin flip ascent to the all-silent local maximum for the highly-repeated stimulus dataset in [41].

We also fit the Tree Hidden Markov model (see *Methods: Fitting the Tree Hidden Markov model*) to the same mildly-repeated stimulus dataset, and then performed - using the stationary distribution of the Markov chain - single spin flip ascent on the corresponding modeled probability landscape. Since the number of free parameters constitutes an important consideration for model comparison, we selected the number of Tree HMM modes so as to match the total number of free parameters as closely as possible to the K-Pairwise MaxEnt model. As shown in Fig 2B, 31 unique local maxima were found after performing single spin flip ascent on the probability landscape modeled by the Tree HMM. Of the 175,002 observed population responses in the data, 99.74% were mapped to the all-silent local maximum, and 0.056% were mapped to the rank-2 local maximum. This result is consistent with that obtained for the K-Pairwise MaxEnt model, indicating that the paucity of true local maxima is an intrinsic feature of the probability landscape rather than an artifact of one specific model.

This exceptionally high percentage of observed states mapped via single spin flip ascent to the all-silent local maximum was even higher when we performed the above procedure on a dataset of 152 ganglion cells responding to a *non-repeated* natural movie stimulus, which we will refer to throughout as Movie #2 (see *Methods: Visual stimulus display*). For this non-repeated stimulus dataset, performing single spin flip ascent on the probability landscape modeled by the Tree HMM identified only two local maxima (see S1 Fig). Moreover, of the 90,001 observed population responses in the data, 99.95% were mapped to the all-silent local maximum (S1 Fig).

To probe the relationship between stimulus repeat structure and prevalence of non-silent local maxima more systematically, we next performed an analysis which allowed us to take the number of stimulus repeats as a parameter. Specifically, two separate experiments were performed in which we presented one of two different movie stimuli: a binary white-noise checkerboard movie, or a natural movie. We will refer to these throughout as Movie #3 and Movie #4, respectively. Both movie stimuli were similarly designed such that unique, non-repeated movie clips were interleaved with presentations of a fixed "target" clip (see *Methods: Parametric Repeat Analysis* for details). A schematic illustrating the experimental setup for the interleaved white-noise checkerboard stimulus (i.e. Movie #3) is shown in Fig 3A.

For each of the two movie stimuli, we then generated "subset movies", which were comprised of a subset of the movie clips appearing in the original full-length movie (see *Methods: Parametric Repeat Analysis*). Different subset movies included a different ratio of the number of repeated target clips to unique movie clips. Note that we held the duration of each subset movie constant, so as to eliminate the potentially confounding effect of different stimulus durations on our sampling of population neural activity. To ensure that there was no dependence on the specific choice of the movie clips included, we also performed multiple independent random permutations to generate the subset movie corresponding with each repeat ratio. For each permutation and for each repeat ratio, we then fit the Tree HMM to the set of RGC population responses observed during the corresponding subset movie. The cross-validated log-likelihood results (used to select the latent dimensionality of the model; see *Methods: Fitting the Tree Hidden Markov Model*) for different repeat ratios are shown in Fig 3B for the interleaved white-noise checkerboard dataset, and in Fig 3E for the interleaved natural movie dataset. For both datasets, there was a general shift toward a larger optimal latent dimensionality as the number of included repeats increased.

After fitting the Tree HMM to the subset of RGC population responses corresponding with each subset movie, we then performed the single spin flip ascent procedure on the modeled probability landscape. Averaged over

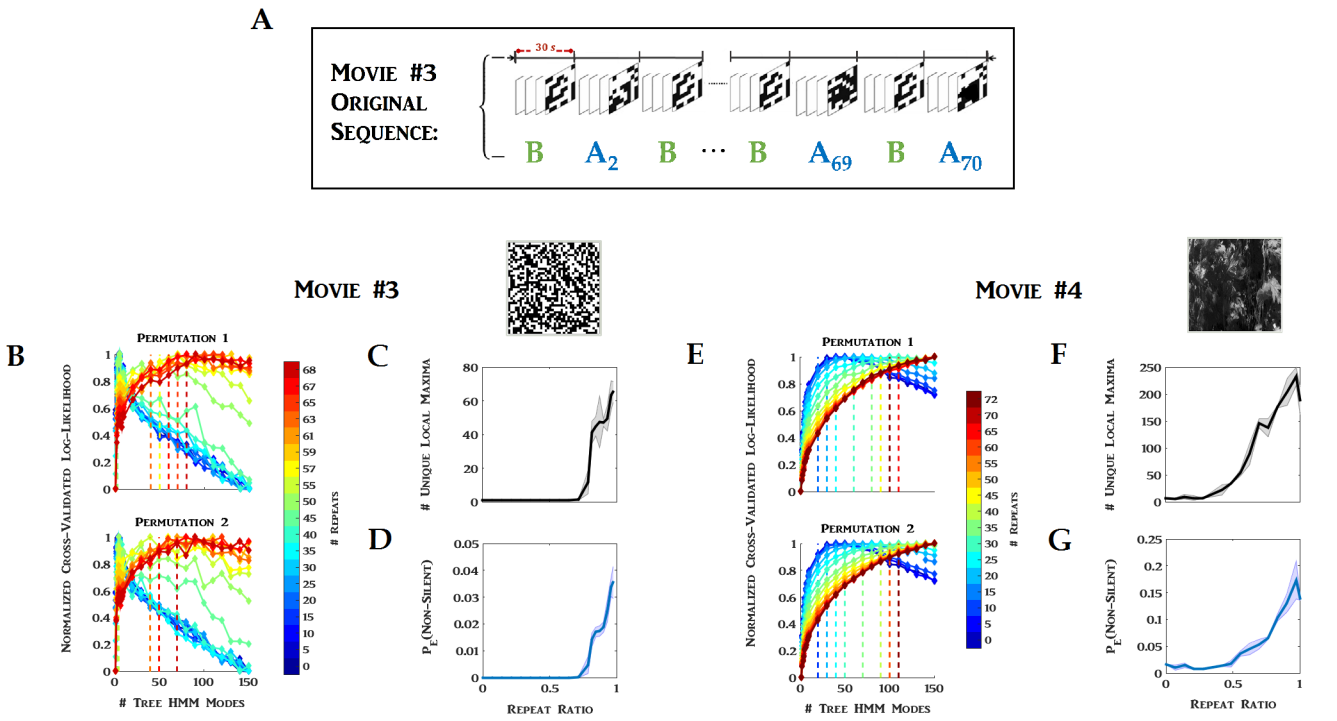

FIG. S3: **Experimental design and results for the parametric repeat analysis. (A)** Cartoon of the full sequence of movie clips comprising the white-noise stimulus (Movie #3). Each movie clip was 30 seconds in duration (red). In the original movie sequence, 69 unique movie clips were interleaved with 69 presentations of a fixed "target" clip (denoted by green B's). The r-th distinct movie clip is denoted by A_r (blue). **(B-D)** Results for a dataset of the responses of $N = 155$ RGCs to Movie #3; an example single frame from this stimulus is shown at top right. **(B)** Normalized cross-validated log-likelihood (CV-LL, y-axis) vs. latent dimensionality of the Tree HMM (x-axis) for each repeat ratio examined. Each color denotes a different number of repeats of the target movie clip which were included in the corresponding subset movie, i.e. equivalently a different repeat ratio (see color key). Results are shown for two of five permutations performed. Dashed lines denote the latent dimensionality at which the corresponding CV-LL curve reaches 90% of of its total height (see *Methods*). **(C)** Total number of unique local maxima (black) found by performing single spin flip ascent on the probability landscape modeled by the Tree HMM vs. repeat ratio. Shown is the average (bold) over all five permutations; shaded error bars denote SEM. **(D)** Proportion of the 105,000 population responses observed during each respective subset movie that were mapped via single spin flip ascent, denoted P_E, to a non-silent local maximum (blue) vs. repeat ratio. Shown is the average (bold) over all five permutations; shaded error bars denote SEM. **(E-G)** Results for a dataset of the responses of $N = 170$ RGCs to the interleaved natural movie stimulus (Movie #4). **(E,F)** Same format as panels (B) and (C), respectively. **(G)** Proportion of the 216,000 joint responses observed during each respective subset movie which were mapped via single spin flip ascent, denoted P_E, to a non-silent local maximum vs. repeat ratio.

all permutations, for both datasets we observed a predominantly monotonic increase in the number of unique local maxima as a function of repeat ratio. Specifically, for the white-noise checkerboard (Movie #3) dataset, we found 1 ± 0 unique local maxima corresponding with the case of 0 included repeats, which increased to 65 ± 6 identified local maxima when 68 repeats were included (Fig 3C). Likewise, for the interleaved natural movie (Movie #4) dataset, increasing the number of included repeats from 0 to 70 corresponded with a drastic increase in the average number of identified local maxima, from 6 ± 0 to 232.5 ± 16.5 (Fig 3F). Moreover, for both datasets, the weight of non-silent local maxima - that is, the percentage of the RGC population responses observed during each respective subset movie which were not mapped via single spin flip ascent to the all-silent local maximum - increased mostly monotonically as a function of the repeat ratio. In particular, for the white-noise checkerboard dataset we found that the weight of non-silent local maxima increased from 0% for the non-repeated case, to $3.54\% \pm 0.59\%$ for the case of 68 repeats (Fig 3D). For the interleaved natural movie dataset, the effect was stronger, with the corresponding values being $1.63\% \pm 0.01\%$ for the non-repeated case, and $17.34\% \pm 3.54\%$ when 70 repeats were included (Fig 3G).

In summary, we have found that the qualitative structure of the RGC population response probability landscape strongly depended on the repeat structure of the stimulus ensemble. More specifically, within the non-

and mildly-repeated stimulus ensemble regime, we found that for both white-noise and natural movie stimuli the modeled response probability landscape was essentially comprised of a single global maximum, the all-silent local maximum, $\vec{0}$. However, when particular stimuli were repeated sufficiently, this resulted in the incorporation of significant non-silent local maxima in the modeled probability landscape. As a check on the validity of these results, we note that the model fitting and spin flip ascent code used here was vetted on multiple analytical test cases; see *Methods: Analytical test cases* and S2 Fig for details.

"Soft" local maxima

In recent past work it was found that for broad, non-repeated stimulus ensembles, the population activity of ganglion cells naturally partitions into a discrete set of "collective modes" [35]. In particular, these partitions were captured by the latent modes formalized in a Tree Hidden Markov model. Moreover, of particular relevance to the present paper, while the exact same input stimulus typically evoked a highly variable set of population response patterns over trials, these disparate response patterns were found to preferentially map to the same mode. This is important, because it suggests that this mode representation could constitute a form of error-correction coding. Moreover, since the authors in [35] identified the modes by fitting the Tree HMM only to the *non-repeated* portion of the three different stimulus datasets they tested (and then calculated the mode reproducibility on the held-out repeat trials), this error-correction capacity was obtained for the non-repeated stimulus regime.

Since these noise-robust collective modes were defined statistically, this leaves open the following question: What is the geometric interpretation of these modes in relation to the response probability landscape? Based on the past maximum entropy work showing a proliferation of non-silent local maxima in the RGC response probability landscape for the *highly-repeated* stimulus regime, we initially expected the collective modes in [35] to correspond geometrically with local maxima. However, we have seen here in the previous section that, for the non- and mildly-repeated stimulus regime, non-silent local maxima are exceedingly rare. This result indicates that collective modes do not necessarily have the geometric structure of a local maximum in the probability landscape. We thus next sought to determine what the empirically-correct geometric interpretation of the collective modes is.

As a first step, we postulated that there is low-dimensional structure in the high-dimensional joint response probability landscape that could potentially be exploited by downstream processing areas. In particular,

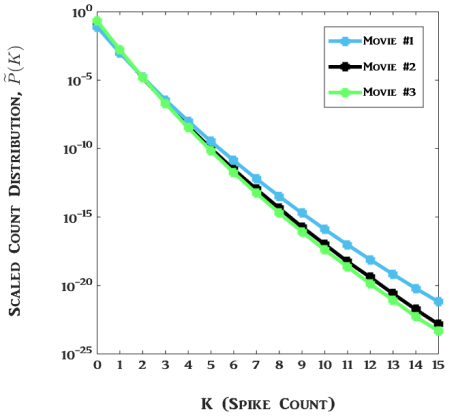

FIG. S4: **Empirical joint response probability vs. spike count.** Shown is the "scaled count distribution", $\widetilde{P}(K)$, as a function of spike count K. The scaled count distribution is defined as the empirical probability of observing a population response with K spikes, normalized by the analytical number of possible joint response patterns with spike count K (see *Methods: Scaled count distribution*). Results are shown for three of the different datasets analyzed in the present paper: Movie #1 (blue; $N = 128$ ganglion cells), Movie #2 (black; $N = 152$ ganglion cells), and Movie #3 (green; $N = 155$ ganglion cells). Note that the y-axis is shown on a log scale.

based on our result in the previous section that sparseness dominates the response probability landscape in the low-repeat stimulus regime (that is, there are exceedingly few local maxima outside of the all-silent global maximum, $\vec{0}$), we speculated that spike count may constitute one important dimension that could be exploited. Note that one conceptual advantage of choosing the spike count as a relevant dimension is that it has an immediate biological interpretation. The importance of spike count in shaping the probability landscape is further suggested by the monotonic decrease in the average normalized empirical probability of a response pattern as a function of the pattern's spike count, a result we consistently observed for all of our datasets (Fig 4).

To explore this idea, we first investigated whether there are identifiable response patterns which are local maxima when we *condition* on spike count. In other words, these states are not necessarily local maxima of the probability landscape, when all possible single spin flips are allowed. To this end, we introduce the following definition:

Definition 1 *A response state* $\vec{\gamma} \in \{0,1\}^N$ *is a* K-soft local maximum *if and only if* $P[\vec{\gamma}] > P[\vec{\sigma}] \;\; \forall\, \vec{\sigma}$ *such that* $w_H(\vec{\sigma}) = w_H(\vec{\gamma}) = K$ *and* $d_H(\vec{\sigma},\vec{\gamma}) = 2$, *where* $w_H(\vec{\sigma}) \equiv \sum_{i=1}^{N} \sigma_i$ *denotes spike count (Hamming weight) and* $d_H(\vec{\sigma},\vec{\gamma}) \equiv \sum_{i=1}^{N} |\sigma_i - \gamma_i|$ *denotes Hamming distance.*

That is, K-soft local maxima are by definition local maxima when we restrict our search to the metric subspace (of the full joint response space, which is the N-

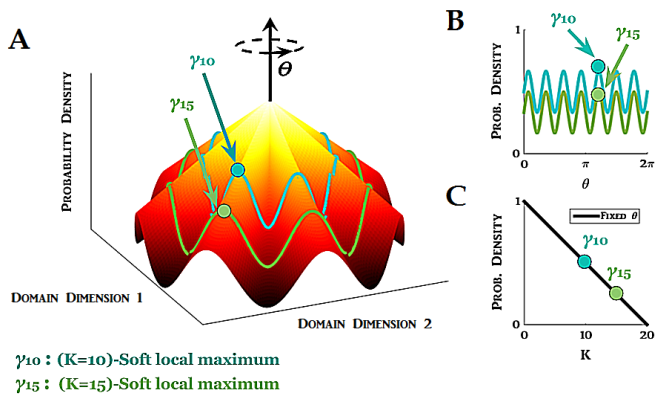

FIG. S5: **Schematic illustrating the concept of K-soft local maxima and ridge points. (A)** Cartoon of a multivariate probability mass function and its "ambient" (i.e. prior to projection) domain. For ease of visualization, the ambient domain is shown in two dimensions, and is represented in polar coordinates. Here we denote the radial coordinate by K (where K assumes positive integer values), and the polar coordinate by θ. If we take the radial coordinate K to correspond with spike count, then the analogue of the K-soft local maxima defined in Def. (1) for this conceptual example are the points which are local maxima of the conditional probability $P(\theta|K)$, i.e. after we project onto a fixed value of K. For illustration, two example K-soft local maxima are shown, denoted by circles and labeled γ_{10} (blue) and γ_{15} (green; see arrows). **(B)** After projecting onto the radial coordinate for the respective values of K, we see that γ_{10} is a local maximum of $P(\theta|K = 10)$ and likewise that γ_{15} is a local maximum of $P(\theta|K = 15)$. **(C)** Illustrated is the projection onto the polar coordinate, for an arbitrary fixed value of θ. Note that since soft local maxima are local maxima of a function defined on a lower-dimensional subspace of the original domain, they are candidates for a type of "ridge point," and as shown in Panel (A) could theoretically form discrete "ridges" which span across different K levels.

dimensional Hamming cube) defined on the set of all possible joint response patterns with fixed spike count K. We will subsequently refer to each such metric subspace as the "K-th spike count level." The geometric intuition behind this definition is illustrated in Fig 5.

In practice, we used the following numerical procedure, which we term "neuron pair relaxation," to find soft local maxima: We start with a population response pattern $\vec{\sigma}(t) \in \{0,1\}^N$ that appears in the data, and attempt to 'flip' a randomly-chosen pair (i, j) of neurons which have opposite responses within time bin t. Specifically, by 'flipping' the response state of neurons i and j, we mean that we set $\sigma_i(t) = \sigma_j(t) \oplus 1$, where $\oplus$ denotes mod 2 addition. Flipping is retained if the resulting pattern has a higher probability than the original pattern (see *Methods: Finding soft local maxima*). This procedure is iterated until none of the $K \cdot (N - K)$ neuron pairs with opposite instantaneous responses can be flipped to increase the probability; the terminating pattern is by definition a soft local maximum. As with our single spin flip ascent code, we vetted the combination of our model fitting and neuron pair relaxation code on analytical test cases to ensure correct implementation (see *Methods: Analytical test cases* and S3 Fig).

To check for the presence (or absence) of soft local maxima in the joint probability landscape, we implemented the above-described neuron pair relaxation procedure for two large datasets. Based on the results of another analysis (see *Methods: Choice of probability model*), which demonstrated that the Tree HMM more accurately captures the empirical $(K = 2)$-soft local maxima than the K-Pairwise MaxEnt model, we chose the Tree HMM as our underlying model of the probability landscape.

We first examined the dataset of 152 ganglion cells responding to Movie #2, which was a 30-min duration non-repeated natural movie stimulus. As shown in Fig 6A, for all spike counts $K > 1$ examined, we found that multiple soft local maxima were always present. Moreover, the number of identified unique soft local maxima monotonically increased with spike count (Fig 6A, grey curve). To investigate how robust the mapping from the dataset of all observed responses with K spikes to the set of unique K-soft local maxima was, for this dataset we performed 100 iterations of our mapping procedure, where each iteration involved a different, independent random sequence of permutations for the neuron pair relaxation algorithm (see *Methods: Finding soft local maxima*). For each value of K, we then computed the mean pairwise overlap ratio between the identified sets of unique K-soft local maxima for each of the 4950 iteration pairs. As seen in Fig 6A (upper left-hand inset), for $1 \leq K \leq 3$, the identity of the mapped set of unique soft local maxima was perfectly conserved. For $4 \leq K \leq 7$, when restricted to iterations for which the cardinality of the mapped set of unique K-soft local maxima was largest, the mean pairwise overlap ratio was also very high ($> 94\%$), indicating a high degree of robustness.

We next examined the responses of 170 ganglion cells responding to the original version of Movie #4, which was a 2.5-hour duration natural movie stimulus in which every 60 s alternated between a unique movie clip and a repeated fixed target clip. As shown in Fig 6B, we similarly found a proliferation of soft local maxima for all examined spike counts. Consistent with the previous results, we observed a monotonic increase in the number of identified unique soft local maxima as a function of spike count; in particular, 227 unique $(K = 10)$-soft local maxima were identified (Fig 6B, grey curve).

"Ridges" are a feature of the probability landscape of neural activity

So far we have restricted our attention to the set of all activity patterns having a fixed number of spikes, K.

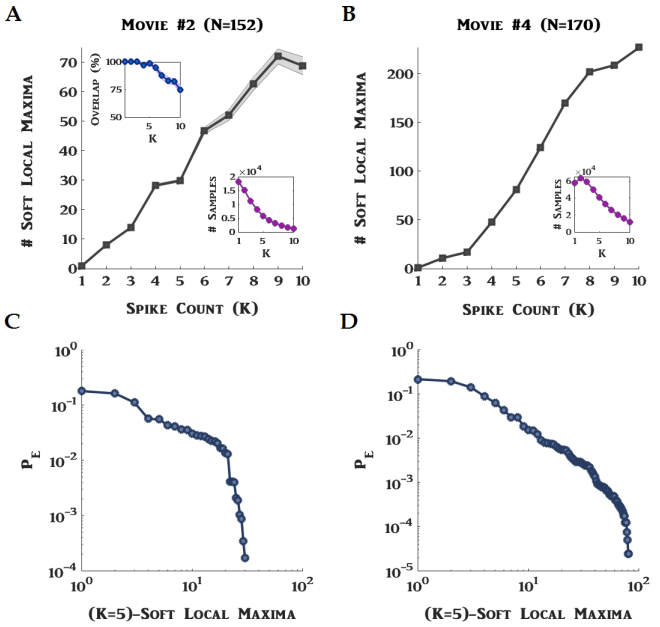

FIG. S6: **Soft local maxima results for two large datasets.** Overall soft local maxima results for **(A,C)** the dataset of $N = 152$ ganglion cells responding to Movie #2, which was a non-repeated natural movie stimulus; and **(B,D)** the dataset of $N = 170$ RGCs responding to the original version of Movie #4 (see *Methods: Visual stimulation*). **(A)** Shown in grey is the mean number of unique K-soft local maxima identified (y-axis) by performing neuron pair relaxation on the response probability landscape modeled by the Tree HMM ($m = 19$ latent modes) for each examined spike count level K, averaged over 100 mapping iterations. Error bars denote one standard deviation over the 100 iterations. *Upper Left Inset*: The mean pairwise overlap ratio (see *Methods: Finding soft local maxima*) for each value of K (blue), averaged over all 4950 iteration pairs. Error bars denote one standard deviation. *Lower Right Inset*: Total number of observed population responses in the data (purple) for each spike count level. **(B)** Same format as in panel (A). For this dataset, one iteration of mapping from observed population responses to the set of unique K-soft local maxima was performed. **(C)** Example results for a specific spike count level K, arbitrarily chosen here to be $K = 5$. Shown is a log-log plot of the proportion of the 5795 observed responses in the data with 5 spikes which were mapped, via neuron pair relaxation, to the corresponding ($K = 5$)-soft local maximum indicated on the x-axis. **(D)** Log-log plot of the proportion of the $40,484$ observed population responses in the data with 5 spikes which were mapped to the corresponding ($K = 5$)-soft local maximum indicated on the x-axis.

Within the disciplines of computer vision and differential geometry, there is a well-studied notion of ridges, which are curves or hypersurfaces composed of so-called "ridge points" [57, 58]. Although multiple definitions for ridge points can be found in the computer vision literature, one popular definition that is considered successful is the *height definition* [58]. This definition stipulates that a necessary condition for a point $\vec{\sigma}$ to be a ridge point of a multivariate function f is that it must be a *generalized maximum*, which conceptually is a local maximum of f when we restrict our search to a subspace of the function's domain. Although generalized maxima are technically only defined for continuous functions on vector space domains, we can see that the K-soft local maxima defined in Definition 1 are a type of discrete analogue, where we have chosen the restricted metric subspace of the domain (which is the Hamming cube in our case) to be the set of all activity patterns with a fixed spike count, K.

Given this potential conceptual connection to ridge points, we thus next sought to investigate whether the soft local maxima found in the previous section are naturally organized *across* spike count levels such that they comprise discrete "ridges." Toward this purpose, we introduce the following definition:

Definition 2 *A soft local maximum $\vec{\gamma}$ is* u-reachable *from another soft local maximum $\vec{\sigma}$ if and only if $w_H(\vec{\gamma}) = w_H(\vec{\sigma}) + 1$, and $\exists$ neuron $i \in$ the set of neurons which have a silent response, $\mathcal{S}(\vec{\sigma})$, such that:*

(i) Performing neuron pair relaxation on $\vec{\sigma}^{(i)}$ results in $\vec{\gamma}$, and

(ii) $P[\vec{\sigma}^{(i)}] < P[\vec{\sigma}]$ (to ensure that $\vec{\sigma}^{(i)}$ is not a local maximum)

where $w_H(\cdot)$ denotes spike count, $\vec{\sigma}^{(i)}$ denotes the joint response pattern which differs from $\vec{\sigma}$ only in neuron i's response, and $\mathcal{S}(\vec{\sigma}) \equiv \{\text{neurons } j \mid \sigma_j = 0\}$.

Intuitively, the motivation behind this definition is to identify the soft local maximum $\vec{\gamma}$ at the $(K+1)$-th spike count level which is 'closest' to the K-soft local maximum of interest (denoted in Def. 2 by $\vec{\sigma}$), while accounting for potential energy barriers. A schematic illustrating the notion of u-reachable as defined in Def. 2 is shown in Fig S7A. This allows us to define a progression of nearest soft local maxima spanning across low to high spike count levels. An analogous notion conducive to exploring soft local maxima organization across high to low spike count levels, which we term d-reachable, is similarly defined:

Definition 3 *A soft local maximum $\vec{\gamma}$ is* d-reachable *from another soft local maximum $\vec{\sigma}$ if and only if $w_H(\vec{\gamma}) = w_H(\vec{\sigma}) - 1$, and $\exists$ neuron $i \in$ the set of neurons which have an active response, $\mathcal{A}(\vec{\sigma})$, such that:*

(i) Performing neuron pair relaxation on $\vec{\sigma}^{(i)}$ results in $\vec{\gamma}$, and

(ii) $P[\vec{\sigma}^{(i)}] > P[\vec{\sigma}]$

where $\mathcal{A}(\vec{\sigma}) \equiv \{\text{neurons } j \mid \sigma_j = 1\}$.

This notion of d-reachable is illustrated in Fig S7B.

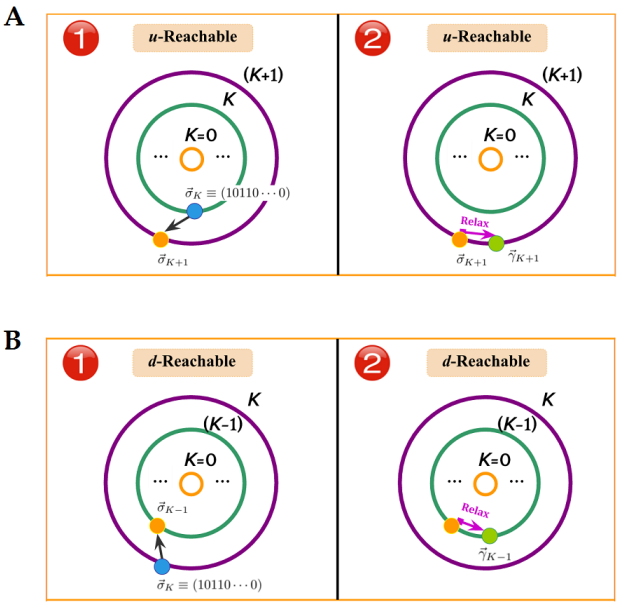

FIG. S7: **Schematic illustrating the notions of *u-reachable* and *d-reachable*. (A)** In this cartoon diagram, $\vec{\sigma}_K$ (blue filled circle) represents a K-soft local maximum, and $\vec{\gamma}_{K+1}$ (green filled circle) represents a $(K+1)$-soft local maximum. The different metric subspaces Ω_K of the full joint response space $\{0,1\}^N$, which are defined as $\Omega_K \equiv \{\vec{\sigma} \in \{0,1\}^N \mid w_H(\vec{\sigma}) = K\}$ and which we also refer to as "spike count levels", are represented by colored concentric rings. The center-most yellow ring corresponds with the 0th spike count level. In this example, $\vec{\gamma}_{K+1}$ is u-reachable from $\vec{\sigma}_K$. This is because: (1) there exists a silent neuron i such that changing neuron i's instantaneous response from $\sigma_i = 0$ to 1 (represented by the black arrow) will (2) result in a joint response pattern $\vec{\sigma}_{K+1}$ (orange filled circle) at the $(K+1)$-th spike count level that will be mapped via neuron pair relaxation (represented by the purple arrow) to $\vec{\gamma}_{K+1}$. **(B)** In this cartoon diagram, the $(K-1)$-soft local maximum $\vec{\gamma}_{K-1}$ (green filled circle) is d-reachable from the K-soft local maximum $\vec{\sigma}_K$ (blue filled circle). This is because: (1) there exists an active neuron i such that changing neuron i's instantaneous response from $\sigma_i = 1$ to 0 (represented by the black arrow) will (2) result in a joint response pattern $\vec{\sigma}_{K-1}$ (orange filled circle) at the $(K-1)$-th spike count level that will be mapped via neuron pair relaxation (purple arrow) to $\vec{\gamma}_{K-1}$.

We visualized examples of soft local maxima organization across low to high spike count levels via a breadth-first search-like (BFS) algorithm, which was used to construct a rooted weighted directed graph (digraph) starting with a given "root" K-soft local maximum (see *Methods: Visualizing ridges* for full details). The nodes of this rooted weighted digraph represent soft local maxima, and a directed edge $\left(\vec{\gamma}_K, \vec{\gamma}_{K+1}\right)$ is drawn if and only if soft local maximum $\vec{\gamma}_{K+1}$ is u-reachable from $\vec{\gamma}_K$. Since the neuron pair relaxation algorithm is not deterministic, for each neuron i in the silent set of the current K-soft local maximum, $\mathcal{S}(\vec{\gamma}_K)$ (of which there are up to $\left(N-K\right)$ choices), we performed 10 iterations of the neuron pair relaxation procedure on $\vec{\gamma}_K^{(i)}$. We then set the weight of each directed edge $\left(\vec{\gamma}_K, \vec{\gamma}_{K+1}\right)$, which we denote by $w(\vec{\gamma}_K, \vec{\gamma}_{K+1})$, as the proportion of the 10 relaxation iterations which resulted in $\vec{\gamma}_{K+1}$. The weight of each edge is represented by edge thickness in the digraph visualizations (Figs 8 and 9). To mitigate effects due to noise, for each unique $(K+1)$-soft local maximum $\vec{\gamma}_{(K+1)}$ reached after performing the above procedure on a given $\vec{\gamma}_K$, we further only included the directed edge $\left(\vec{\sigma}_K, \vec{\gamma}_{K+1}\right)$ if the corresponding edge weight $w(\vec{\gamma}_K, \vec{\gamma}_{K+1})$ exceeded a reliability threshold Θ. In practice, we used $\Theta = 0.3$.

Via this approach, we thus algorithmically define a discrete ridge as the set of soft local maxima which are organized across low to high spike count levels such that each respective pair of K- and $(K+1)$-soft local maxima is 'connected', in the sense of being u-reachable. We show example rooted weighted digraphs obtained for the dataset of 152 ganglion cells responding to the non-repeated natural movie (Movie #2), for the choice of starting spike count level $K = 4$ and terminating spike count level $K = 17$ (Figs 8 and 9). In most examples, a u-reachable soft local maximum had one active neuron added; i.e. that neuron was silent at spike count level K and spiking at level $K+1$. Many digraphs exhibited one or a few neurons added at each u-reachable step; these digraphs formed simple ridges.

Note that in theory, it is possible that a discrete ridge descending from the global maximum (which occurs for this dataset at the 0th spike count level) could terminate at a lower spike count level K than the one we arbitrarily choose our numerical procedure to explore out to; an example occurrence of this scenario is shown in Fig 8C. Further, it is theoretically possible for a ridge to branch into multiple ridges. Consequently, each output rooted digraph could potentially contain more than one ridge. To aid in visualizing distinct ridges, we further assigned an x- and y-coordinate to each node in the output rooted digraphs, with the intuition behind the x-coordinate assignment being to quantify whether or not there is a 'jump' between ridges based on the amount of overlap between the different soft local maxima (see *Methods: Visualizing ridges*). Since the u-reachable qualification alone does not capture the scenarios of early ridge termination and branching, we then extended our algorithmic definition of a distinct ridge by further qualifying that it must be comprised of soft local maxima which share the same x-coordinate, as formulated in Eq. 11.

Connection between Ridges and Tree HMM Modes

For each soft local maximum $\vec{\gamma}$ appearing as a node in the rooted digraph visualizations, we also computed the maximum *a posteriori* (MAP) estimate of its associated Tree HMM mode α, given by Eq. 14 (see *Methods: Vi-*

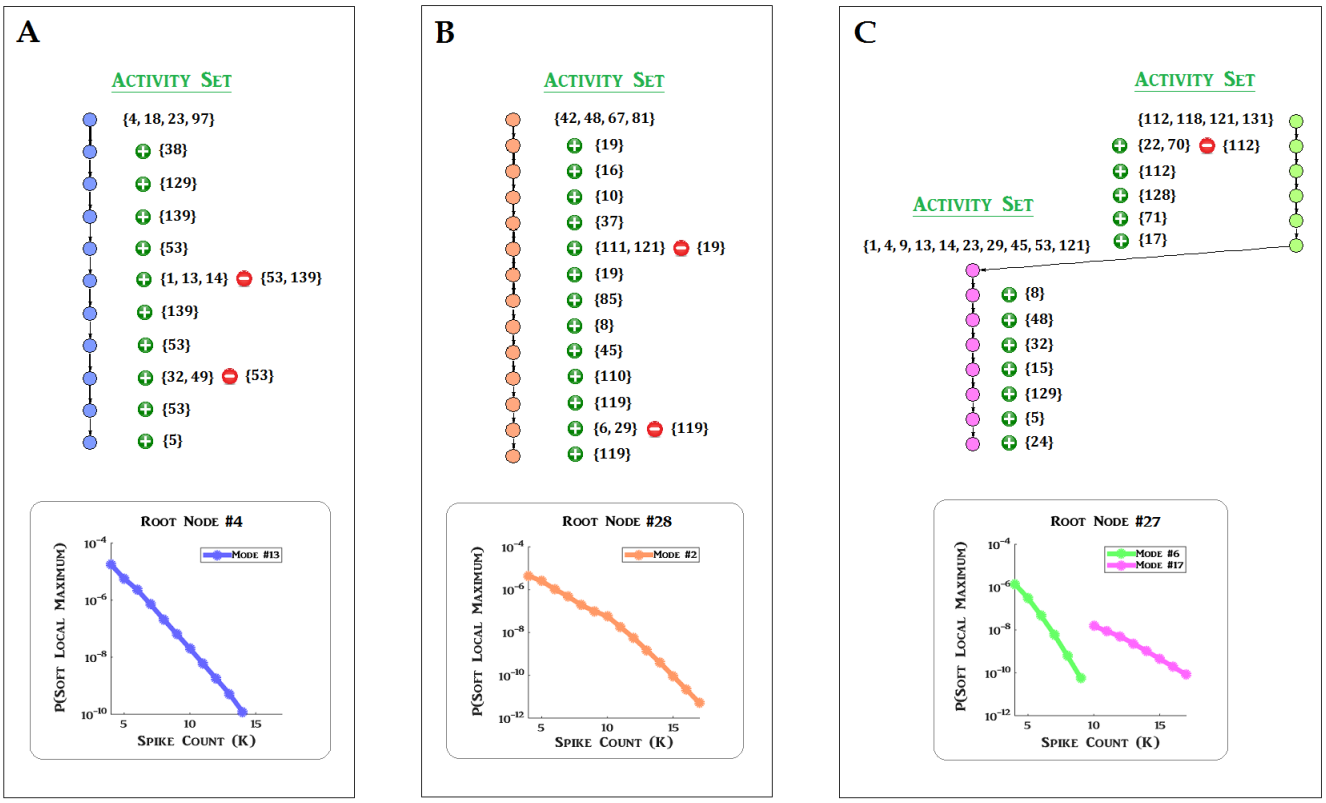

FIG. S8: **Ridges of Type I.** Example Type I (see main text for definition) rooted weighted digraphs obtained for the dataset of $N = 152$ RGCs responding to the non-repeated natural movie. **(A)** Output weighted rooted digraph obtained via our BFS procedure, described in the main text, when we took the unique $(K = 4)$-soft local maximum with rank 4 (i.e. the 4th most common soft local maximum to which all observed population responses with four spikes were mapped) as the input root node. Labeled next to each node is the "activity set" of the corresponding soft local maximum $\vec{\gamma}$, defined as $\mathcal{A}(\vec{\gamma}) \equiv \{i \,|\gamma_i = 1\}$. For ease of visualization, the activity set of the starting node is written out in full. The activity set of each subsequent $(K+1)$-soft local maximum, $\vec{\gamma}_{K+1}$, is then denoted by the neuron(s) which is added to (green plus symbol) or removed (red minus symbol) from the previous K-soft local maximum's activity set to achieve $\mathcal{A}(\vec{\gamma}_{K+1})$. Activity sets are written out in full in S4 Fig. Node color denotes the MAP estimate of the associated latent Tree HMM mode. Edge weights are represented by edge thickness; thicker edges denote larger weights. There was one distinct ridge for this example; all soft local maxima were associated with the same MAP estimate, mode #2 (pink). Below: Probability, as modeled by the Tree HMM, of each depicted soft local maximum vs. its spike count, K. **(B)** Rooted digraph obtained when the $(K = 4)$-soft local maximum with rank 28 was taken as the input root node. Notation is the same as in panel (A). There was one distinct ridge for this example; all soft local maxima were associated with the same MAP estimate, mode #13 (light blue). **(C)** Rooted digraph obtained when the $(K = 4)$-soft local maximum with rank 27 was taken as the input root node. There were two distinct ridges. All soft local maxima comprising the first ridge were associated with mode #6 (light green); those comprising the second ridge were associated with mode #17 (pink).

sualizing ridges). To visualize this information, we then assigned a color to each node in the output rooted digraphs, which uniquely identifies the MAP estimate of the associated Tree HMM mode.

Overall, there was a strong tendency for each distinct ridge to be comprised of soft local maxima which shared the same MAP estimate for the associated Tree HMM collective mode (Fig 8). To broadly characterize the digraph visualization results in terms of the degree of correspondence between ridges and collective modes, we classified each unique ridge as belonging to one of two disjoint categories, which are defined in Table I. Conceptually, a "Type I" classification denotes that the given ridge has a 'perfect' correspondence with one of the latent Tree HMM modes. Note that we chose to classify on the basis of ridges, rather than on the basis of the rooted digraph examples, to avoid potentially double-counting ridges when computing global statistics. Applying this classification scheme for the dataset of 152 RGCs responding to Movie #2, we observed that 13 of the 17 total identified unique ridges, i.e. 76.5%, were Type I. Comparably, for the dataset of 170 RGCs

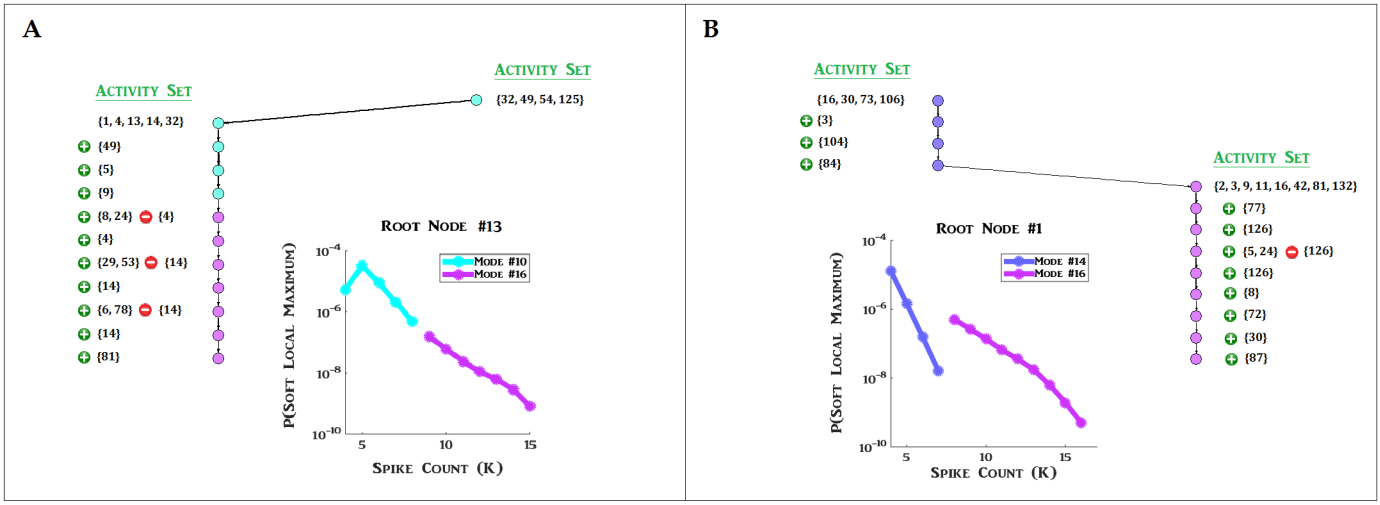

FIG. S9: **Ridges of Type II.** Example Type II (see main text for definition) rooted digraphs obtained for the dataset of $N = 152$ RGCs responding to the non-repeated natural movie. **(A)** Rooted weighted digraph obtained when the $(K = 4)$-soft local maximum with rank 13 was taken as the input root node. Notation is the same as in Fig 8. There was one distinct ridge; however, the soft local maxima comprising this ridge were associated with two possible MAP estimates: mode #10 (cyan) and mode #16 (violet). **(B)** Rooted digraph obtained when the $(K = 4)$-soft local maximum with rank 1 was taken as the input root node. There were two distinct ridges. The first, which was comprised of soft local maxima which were all associated with the same MAP estimate (mode #14, blue), was Type I. Soft local maxima comprising the second ridge were associated with mode #16 (violet). Since mode #16 corresponded with three distinct ridges, two of which are shown here in panels (A) and (B), this second ridge was classified as Type II.

responding to the original version of Movie #4, we found that 32 of the 44 total identified unique ridges, i.e. 73%, were Type I. Hence for both a non-repeated as well as a moderately-repeated stimulus ensemble, we found that there was a substantial correspondence between the statistically-defined Tree HMM modes and the geometrically-defined discrete ridges.

Connection between Ridges and "Neuron Communities"

As shown in Fig 9, even in the few cases where there was not a perfect one-to-one bijective mapping between a given geometric ridge and a Tree HMM mode, we consistently found that each identified ridge corresponded with a specific group of active neurons that we will here term a "neuron community." In particular, for each distinct ridge, a group of neurons in the population could be identified such that every soft local maximum $\vec{\gamma}$ comprising the ridge had an active set of neurons, defined as $\mathcal{A}(\vec{\gamma}) \equiv \{\text{neurons } i \mid \gamma_i = 1\}$, which was a subset of this identifiable group. Neurons outside of this identifiable "neuron community" were always silent for all soft local maxima comprising the associated ridge.

To investigate how the full population clusters into these ridge-associated neuron communities, for each neuron i (where $1 \leq i \leq N$) we recorded all of the neuron communities it was a member of based on the compilation of the rooted digraph results. We then visualized this information in the form of an undirected graph, where each node in the graph represents a neuron in the population, and an edge (i, j) is present if and only if neurons i and j are members of at least one common neuron community (see *Methods: Visualizing neuron communities*). By definition, the maximal cliques of this graph, which we refer to as the "ridge union graph" for the population, correspond with the ridge-associated neuron communities. The ridge union graph for the dataset of 152 ganglion cells responding to the non-repeated natural movie is shown in Fig 10. As seen in Fig 10, mixed membership, i.e. the case of a neuron having membership in more than one neuron community, was prevalent. More specifically, averaged over the 83 retinal ganglion cells which were assigned to at least one neuron community, a given RGC belonged to a mean ($\pm$ SEM) of 2.5 ± 0.18 neuron communities. This high degree of mixed membership may be a signature of a combinatorial neural population code.

Finally, we sought to ascertain that correlations in the data were necessary to give rise to the soft local maxima and discrete ridge structures observed here, as opposed to these features trivially arising from the structure of the underlying Tree HMM probability model. Due to the large population sizes of the datasets we analyzed, it is intractable to empirically determine the joint response p.m.f., and consequently intractable to directly investigate the presence of these geometric features in the empirical probability landscape. To control for the

TABLE I: **Ridge Classification Categories**

Ridge Type	Description
I	The ridge corresponds with exactly one Tree HMM mode, and this mode moreover uniquely corresponds with this ridge. That is, there is a bijective mapping.
II	There is not a bijective mapping between this ridge and one of the Tree HMM modes. E.g. the ridge is associated with more than one Tree HMM mode, and/or the mapped Tree HMM mode is associated with more than one ridge.

possibility that these features are solely attributable to our specific choice of the underlying probability model, we compared the results obtained for the true dataset of 152 RGCs responding to Movie #2 with those obtained by fitting the Tree HMM to two manipulated cases in which only limited structure was preserved (Fig 11). Details of the construction of both controls are reported in *Methods: Shuffled controls*. In brief, both controls were constructed so as to completely eliminate the signal and noise correlations in the original data; one of the two controls retained heterogeneous firing with the ganglion cell population, while the other control imposed homogeneous firing.

As seen in Fig 11A, there was a categorical difference between the cross-validated log-likelihood (CV-LL) curves obtained when the Tree HMM was fit to the original data, versus when the Tree HMM was fit to the control data. Whereas the CV-LL curve for the original data exhibited a well-defined peak at 19 modes, when fit to each control dataset, the CV-LL curve instead achieved its maximum at only one mode. The soft local maxima results were likewise substantially different for the original data versus the controls (Fig 11B). In particular, there was a monotonic increase in the proliferation of soft local maxima at all spike count levels for the original data, with 74 unique soft local maxima identified at spike count level $K = 10$.

In contrast, the modeled probability landscapes for both controls exhibited a constant and low number of unique soft local maxima across all spike count levels. Specifically, for the first control, in which the firing rates of each neuron in the original data were preserved and thus heterogeneous over the population, one unique soft local maximum was present at each spike count level. As expected analytically for an independent neuron population with heterogeneous firing rates, we confirmed that the single soft local maximum identified via our procedures at each spike count level K corresponded with the K ganglion cells which had the highest firing rates (S6 Fig). For the second control, in which each neuron was moreover assigned the same firing rate, no soft local maxima were found for any of the tested spike count levels. This is consistent with the analytical result and expectation for an independent population of homogeneous neurons, as the symmetry in the homogeneous case implies that the probability landscape depends only on K (and thus all joint responses within the same spike count level have identical probability). Thus, the soft local maxima results, and by extension the ridge results, obtained via our approach for the original data are non-trivially dependent on the empirical correlation structure of the measured neural activity.

Discussion

It is well established that the manner in which local neural circuits encode information generally involves the cooperation of many neurons [8], and that noise is prevalent throughout the nervous system [44–47]. Partitioning noisy output patterns into fault-tolerant clusters corresponding with "codewords" is a long-standing idea which has been utilized in traditional error-correcting codes in communications engineering [51]. However, whether this idea is actually realized in neuronal population codes is a question which has only very recently begun to be empirically investigated [17, 35, 41, 63]. In particular, two independently-conducted studies employing activity (i.e. unsupervised) models to investigate RGC population response data have revealed support for this idea, but via different definitions of codewords [35, 41]. Specifically, the results in Ref. [41] suggested that geometrically-defined local maxima in the joint response probability landscape (which was modeled via a MaxEnt model) could potentially correspond with codewords, whereas the results in Ref. [35] suggested that potential candidates for codewords are captured by the statistically-defined latent modes of the fit Tree hidden Markov model.

In this work, we first sought to investigate how these two different definitions of RGC population codewords are related. We have shown that the qualitative structure of the RGC joint response probability landscape, which is fundamentally shaped by correlations among ganglion cells, depends on the stimulus ensemble statistics. In particular, we have found that for both a natural and artificial movie stimulus, the number of non-negligible local maxima comprising the modeled probability land-

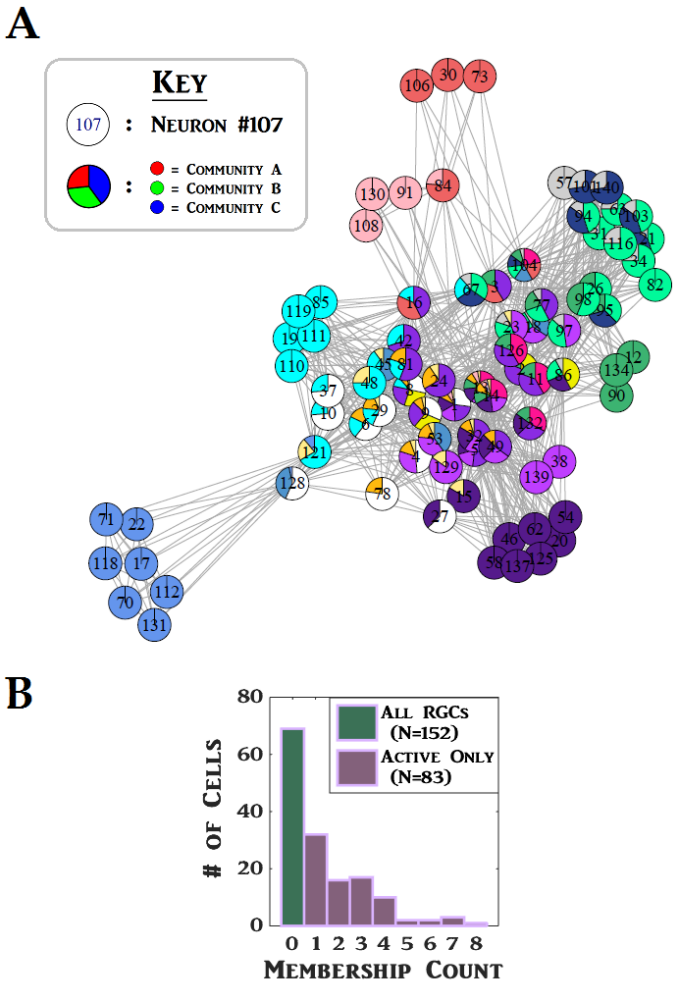

FIG. S10: **Neuron Communities. (A)** Shown is the "ridge union graph" for the dataset of $N = 152$ RGCs responding to the non-repeated natural movie stimulus, i.e. Movie #2. Nodes in this undirected graph represent individual RGC neurons in the population, where each inscribed number indicates the index of the corresponding neuron. The presence of an edge (i, j) denotes that neurons i and j share at least one common ridge-associated neuron community (see *Methods: Visualizing neuron communities*). Each color uniquely identifies one of the 17 distinct ridge-associated neuron communities for this dataset. Neurons which exhibited mixed membership - i.e. were members of more than one neuron community - are depicted with "pie" nodes; in this case, each constituent color denotes the respective neuron community (see Key). **(B)** Histogram of the breakdown of the different degrees of mixed membership across the population (for the same dataset as in panel A). Represented on the x-axis is the number of ridge-associated communities a given ganglion cell was a member of, which we term the "membership count." Green denotes the same as purple, except that it also includes the 69 RGCs which were not a part of any community. The purple histogram only includes the 83 RGCs which were a member of at least one community.

scape is strongly influenced by the number of stimulus repeats (Fig 3). Whereas a sufficiently high-repeat stimulus regime corresponded with a response probability landscape characterized by a proliferation of local maxima, consistent with past results [41], in the low-repeat stimulus regime, we found instead that only a negligible number of activity states were mapped onto non-silent local maxima (Figs 2 and 3). Yet for the natural movie stimulus in the non-repeated regime, there were still many Tree HMM modes (Fig 3E; [35]). Thus, at least in the low-repeat stimulus regime, the geometrically-defined local maxima examined in Ref. [41] and the statistically-defined Tree HMM modes introduced in Ref. [35] are not the same. We then sought to determine the geometric correlate of the Tree HMM modes in the non-repeated stimulus regime. Using a new method of analysis, we showed that these collective modes correspond closely with algorithmically-defined discrete "ridges" in the response probability landscape (Figs 8 and 9). Moreover, we found that each ridge also corresponded with an active neuron "community", a notion from computer science [55].

This organization of the RGC population into a community structure with a high degree of mixed membership (Fig 10) is reminiscent of the concept of "cell assemblies" hypothesized over sixty years ago by Donald Hebb [64]. In Hebb's view, individual neurons did not work in isolation, but rather were suggested to participate in multiple different cell assemblies [10, 64]. The first four defining properties of Hebb's original cell assembly concept are (paraphrased; see [8, 64] for details): (*i*) overlapping set coding of information items; (*ii*) sparse coding; (*iii*) dynamic construction and reconstruction; and (*iv*) dynamic persistence. The neuron community results found here are highly consistent with properties (*i*) and (*ii*). In particular, we found that communities are overlapping sets of neurons: for the non-repeated natural movie dataset, a given ganglion cell was a member of on average 2.5 ± 0.18 different communities, and a given community overlapped with on average 10.7 ± 0.59 other communities. Moreover, any individual community contained a small subset (on average 12.1 ± 1.2 RGCs) of the 152 total neurons in the population, consistent with the sparse coding property. We cannot comment on how our ridge-associated neuron community results relate to Hebb's properties (*iii*) and (*iv*), since these are dynamic properties, and the geometric ridges and consequently communities are features which were extracted from the modeled *static* probability landscape (see *Methods*).

The neuron community results obtained here are also consistent with the fifth and final defining property of Hebb's original cell assembly concept: the dynamic completion property, which stipulates that activation of a large enough subset of a cell assembly results in activation of the complete cell assembly [64]. Although intrinsic dynamics are not applicable in our case, there is a

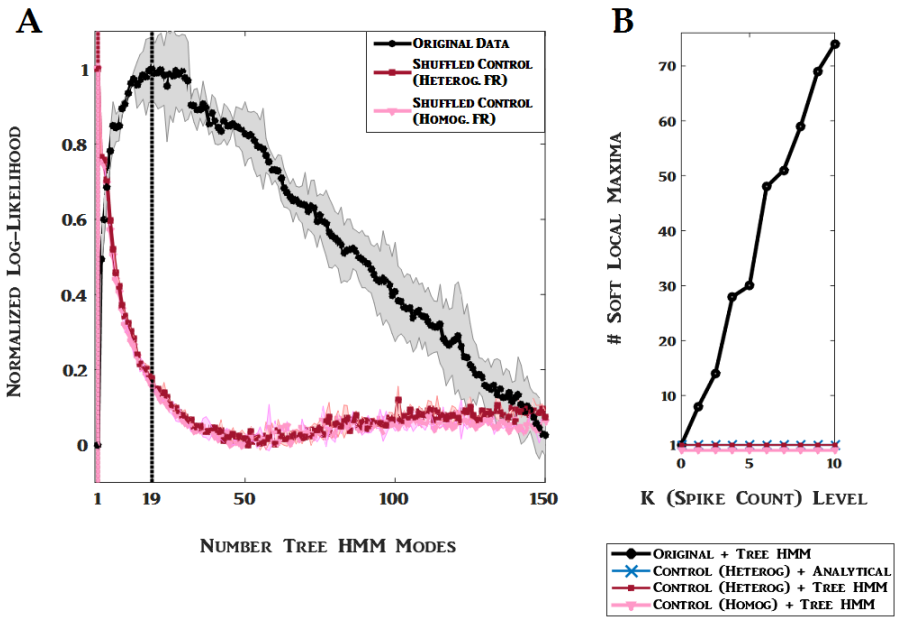

FIG. S11: **Control results. (A)** Normalized cross-validated log-likelihood (CV-LL) as a function of the number of Tree HMM modes for: the original dataset of 152 RGC responses to Movie #2 (black), the heterogeneous firing rate shuffled control (red), and the homogeneous firing rate shuffled control (pink). See *Methods: Shuffled controls* for details of the construction of both controls. Shown in bold is the mean normalized CV-LL over all n cross-validation folds (in practice taken to be $n = 2$). Shaded error bars denote SEM over cross-validation folds. Each colored dashed line indicates the optimal latent dimensionality for the corresponding dataset (see Key for colors). **(B)** Shown is the number of unique K-soft local maxima identified via performing our neuron pair relaxation procedure, as a function of spike count K. Results are shown for four cases: the original dataset with the Tree HMM chosen as the underlying probability model (black circles); the heterogeneous firing rate shuffled control with the analytical independent model as the underlying probability model (blue crosses); the heterogeneous firing rate shuffled control with the fit Tree HMM as the underlying probability model (red squares); and the homogeneous firing rate shuffled control with the Tree HMM as the underlying probability model (pink triangles).

strong parallel between property (v) and our observation that population responses which had instantaneous activation of a large enough subset of the neuron community were typically mapped to the same Tree HMM mode. It was shown in Ref. [35] that the Tree HMM modes were robustly reproducible across stimulus repeats, which combined with the present work suggests that communities are likewise fault-tolerant coding features. In combination, the three properties of overlapping set coding, sparse coding, and robust completion were previously shown to offer theoretical advantages to overcome the deficiencies of the single-neuron doctrine [65–68]. In particular, the development of these ideas in the older literature was essentially motivated by considerations of the trade-off between rate (i.e. the number of information items which could be represented by the assemblies) and error-correction performance; this fundamental trade-off has been rigorously formalized in communications engineering [52].

Within the more modern body of literature in the field, there have been many different proposals for how to study population neural codes. One particularly formative approach has been to apply information-theoretic techniques [2, 24, 62, 69, 70]. In past work which applied such information-theoretic methods to quantify combinatorial coding in the salamander and guinea pig retina, it was found that whereas synchronous spiking was mostly redundant, combinations of spiking and silence were generally synergistic [62]. In the present work, we found that ganglion cells in the population cluster into groups (which we call "neuron communities") that are specific to each ridge identified in the probability landscape. Within a given ridge, population responses exhibit active neurons that are members of this identifiable group, combined with silence from all neurons outside of this group. This result is consistent with the previous synergy from silence results [62], and can be thought of as a generalization of the previous result to populations of $N > 100$ neurons. In particular, we found that neural activity patterns having different numbers of active neurons but the same silent set were mapped to the same ridge. Thus, the number of active neurons within a given community embodies a form of error-correction, while the identity of the silent neurons is, in some sense, more important in defining the population codeword.

Based on a similar motivation to understand the role

of network noise on the RGC population code, another recent approach has been to investigate the possibility of clustering in the response space using a *semantic* distance metric [17]. Using a hierarchical agglomerative clustering algorithm on the matrices formed by computing the semantic distance between each pair of observed joint responses for populations of 20 RGCs, it was found that the resulting ordered matrices exhibited an almost perfect block structure [17]. This result was shown to be non-trivial and indicates that, empirically, the 20-bit RGC population responses cluster in the response space based on semantic similarity. Unfortunately, we cannot directly compare the ridge-associated communities found here with the previously-reported semantic clusters, since the approach in Ref. [17] is a supervised approach that requires a high-repeat stimulus regime, whereas our analyses in the present work were applied to the non-repeated stimulus regime. Given our finding that the RGC joint response probability landscape is proliferated by local maxima for the high-repeat regime, we believe that the previously-reported semantic clusters largely correspond geometrically with local maxima.

However, we speculate that the probability landscape in the vicinity of such local peaks still has an extended, ridge-type geometry for the highly-repeated dataset analyzed in Ref. [17]. In this case, the finding in Ref. [17] that simple linear and bilinear distance metrics do not correspond with semantic distance is also consistent with our findings, as we would not expect simple linear or bilinear clustering routines to succeed well in capturing this type of asymmetric syntactic structure. Finally, our community results are also highly consistent with the finding in Ref. [17] that population responses belonging to the same semantic cluster exhibited some shared structure in terms of having a common subset of neurons which were always active, and other neurons which were always silent. Our results make a further connection between clustering structure within the functional network of RGCs, and geometric features (specifically ridges) in the joint response probability landscape.

Another related, recently-proposed approach has been to use a perturbed maximum entropy framework to characterize the organization of RGC population responses in the response space [63]. By combining a mean-field approximation and the message passing algorithm [72], the authors in Ref. [63] were able to estimate the number of population responses at a given Hamming distance from a chosen reference population activity pattern. Conceptually, introducing a strong perturbative coupling field that is proportional to a given neural activity state corresponds with a probability distribution that is an approximate delta function around the chosen reference state. Then, reducing the perturbation allows the probability to relax into larger and larger neighborhoods around the reference state. Using this approach for both Hopfield networks and retinal network data, it was found that the response space was partitioned into multiple clusters; moreover, for the retinal network data, there was greater clustering around local maxima than around non-local maxima [63]. It is important to note that the retina network dataset analyzed in Ref. [63] was the same as in Refs. [20, 41]. Consequently, these results pertain to the highly-repeated stimulus regime, i.e. the regime in which local maxima proliferate the probability landscape. To our knowledge, our work here is the first to extend similar clustering results for RGC population responses to the non/mildly-repeated stimulus regime.

When studying neural population codes, it is important to make the distinction between the symbolic structure of the code, versus the mapping between inputs and output responses. Here we have addressed only the former of these two issues. This leaves open the question: Do the ridges and/or communities have any meaningful relationship to the stimulus ensemble, and if so, what stimulus features do they encode? For the low-repeat stimulus regime, we have established numerically that the Tree HMM modes appear to correspond closely with geometric ridges in the joint response probability landscape, and moreover with neuron communities. Thus, by extension, we would expect that the ridges and neuron communities encode similar stimulus features as the corresponding Tree HMM modes. Previous work has shown that coding symbols based on multi-neuron responses, including the Tree HMM modes, encode stimulus features which are distinct from those encoded by single neuron responses [35, 62, 71]. Via estimating each Tree HMM mode's receptive field (RF) as the mode-triggered stimulus average, it was found in Ref. [35] that approximately half of the mode RFs resembled the intersection of many individual ganglion cells' RFs. This property represents a straightforward mechanism by which the retinal population code can achieve greater spatial acuity than individual cells. Similar intersection features have been reported to be encoded by the simultaneous spiking patterns of groups of 2-4 RGCs [71]. Additionally, other Tree HMM mode RFs were found to resemble either the union of the underlying individual cells' RFs, or intriguingly, an oriented dipole profile [35]. These properties may be useful for enabling a form of position invariance, and detection of oriented contrast edges, respectively. In all of these cases, the visual features represented by Tree HMM modes were different from those represented by individual ganglion cells. Thus, the set of all Tree HMM modes appears to constitute a new "vocabulary" with which to represent the visual world.

The fact that the geometry of the probability landscape of ganglion cell population activity is qualitatively different in the high-repeat versus low-repeat stimulus regimes raises an obvious question: which regime corresponds to the visual stimuli falling on the eye during real behavior? Answering this question depends on understanding on a more fundamental level how a "stimu-

lus ensemble" is defined from the viewpoint of the retina and downstream brain areas. In particular, there are two component considerations: (1) the duration of sampling of the stimulus space, and (2) whether readout processing is context-specific. Regarding (1), in the real animal, readout circuits can potentially accumulate sampling over long time periods. In the case that sampling occurs over the entire lifetime of the organism, then the corresponding stimulus ensemble will be over all possible natural visual stimuli in the animal's environment, which has been shown to have conserved second-order statistics [73]. Consequently, for the retinal code in this limit, the stimulus ensemble would lack repeated features. Thus, for the case in which sampling occurs over the lifetime of the animal, we would expect that the RGC joint response probability landscape more closely resembles the ridge (i.e. non-repeated stimulus) regime studied here.

However, regarding (2), past work has shown that the neural code can be multiplexed, such that there is not necessarily one single interpretation extracted by downstream processing areas, but instead multiple context-specific interpretations which serve different purposes [74–77]. It is possible that some downstream circuits which read out the ganglion cell population code could be gated by contextual signals, i.e. that these readout circuits only sample from the ganglion cell population under particular contexts. For example, an animal being in different environmental locations could constitute different contexts, signaled by feedback from the hippocampus; another example is running versus stationary movement, which is known to induce contextual modulation in V1 responses [78]. These types of contextual signals could therefore restrict sampling to subsets of the visual environment in which some visual features repeat frequently. In this case, the probability landscape of ganglion cell population activity would have at least some local peaks that downstream circuits might benefit from identifying.

Within the scope of the full information processing system of the animal, these two regimes of sampling over the animal's life span versus contextually-modulated sampling are not mutually exclusive. Indeed, we would suggest that both regimes are actively realized in the full system. Accordingly, we do not interpret our results as necessarily implying that true local maxima are impracticable candidates for error-robust "codewords". Rather, we would speculate that a more likely alternative is that the retinal population code contains at least two *types* of error-robust "codewords": e.g. ridges may constitute one type which encode more general stimulus features such as classes of related stimuli, whereas local maxima may constitute another type which encode more specific, individual stimulus features that occur in a given context. The question of how these two different types of retinal population "codewords" are decoded biologically by downstream processing areas is, of course, an open and important question for future work.

Methods

Ethics statement

This study was performed in strict accordance with the recommendations in the Guide for the Care and Use of Laboratory Animals of the National Institutes of Health. The protocol was approved by the Institutional Animal Care and Use Committee (IACUC) of Princeton University (Protocol 1828).

Experimental procedures

Electrophysiology

We analyzed recordings from larval tiger salamander (*Ambystoma tigrinum*) retinal ganglion cells responding to either naturalistic or white noise checkerboard movie clips. In brief, animals were euthanized according to institutional animal care standards. For the white-noise checkerboard experiment, eyes were removed in darkness and the retina separated from the pigment epithelium. The natural movie experiments were conducted with the pigment epithelium intact and surgery performed in dim light. The isolated retina was then transferred as quickly as possible into oxygenated Ringer's medium at room temperature, in order to optimize the long-term stability of recordings. Tissue was flattened and attached to a dialysis membrane using polylysine. The retina was then lowered with the ganglion cell side against a custom 252-electrode multi-electrode array. Raw voltage traces were digitized and stored for offline analysis using a 252-channel preamplifier (MultiChannel Systems, Germany). Offline spike sorting was performed using custom software. Detailed recording and analysis methods, including the spike-sorting algorithm, are described in [20]. For the three natural movie experiments, 128, 152, and 170 neurons passed the standard tests for waveform stability and lack of refractory period violations, respectively. For the binary white noise checkerboard experiment, 155 neurons passed the standard tests.

Visual stimulus display

Stimuli were projected onto the array from a CRT monitor at a frame rate of 60 Hz, and gamma corrected for the display. We presented two stimulus classes, in a total of four different experiments: a natural movie (three experiments), and a binary white noise checkerboard (one experiment). The natural movie consisted of a 7-min gray scale recording of leaves and branches blowing in the wind. We conducted the natural movie experiments with three different designs: (1) in the first, which we refer to throughout "Movie #1", the full 7-min movie

was looped (i.e. repeated) ten times, for a total movie length of approximately 70 min; (2) in the second, which we refer to as "Movie #2", the movie was looped five times, but with a different pixel centered on the recording area for each repeat; (3) in the third, which we refer to as "Movie #4", we interleaved 73 unique movie clips, each 60 s in duration, with 73 repeats of a fixed 60-s "target" movie clip. Design (2) was constructed this way to provide a non-repeated stimulus: since the patch of retina recorded by the multi-electrode array subtended only a small portion of the stimulus, the retinal input was effectively non-repeated over the full recording session. The binary white noise stimulus, which we refer to throughout as "Movie #3", consisted of a 40×40 array of $58\,\mu m$ squares whose light intensity was randomly selected to be either bright or dark every 30 ms. This flickering checkerborad stimulus was formatted into 30 s periods that alternated between non-repeated (unique) stimulation and a repeated 30-s "target" clip. A total of 69 unique 30-s clips, and 69 repeats of the same 30-s target clip, were presented.

Data preparation

For each dataset, we discretized spike trains into 20 ms time bins, as in previous work [17, 22, 35, 41]. This produces a sequence of binary population response vectors, which we denote as $\vec{\sigma}(t) \equiv (\sigma_1(t), \cdots, \sigma_i(t), \cdots \sigma_N(t)) \in \{0,1\}^N$, where $i = 1, \cdots, N$ labels the neuron identity and t the time bin. We set $\sigma_i(t) = 1$ whenever neuron i fired at least one spike within bin t, and 0 otherwise.

Fitting the K-Pairwise Maximum Entropy model

The analytical form of the K-Pairwise maximum entropy distribution is [41]:

$$P^{(2,K)}\left(\vec{\sigma}\right) = \frac{1}{Z}\, e^{-\mathcal{H}(\vec{\sigma})} \quad \text{where} \quad Z = \sum_{\vec{\sigma}\in\{0,1\}^N} e^{-\mathcal{H}(\vec{\sigma})} \tag{1}$$

and where the "energy function" $\mathcal{H}$ is given by:

$$\mathcal{H}\left(\vec{\sigma}\right) = \sum_{i=1}^{N} h_i \sigma_i + \frac{1}{2}\sum_{i\neq j} J_{ij}\sigma_i\sigma_j + \sum_{K=0}^{N} \lambda_K \delta_{K,w_H(\vec{\sigma})} \tag{2}$$

where $\delta_{a,b}$ is the Kronecker delta (i.e. $\delta_{a,b} = 1$ if $a = b$ and 0 otherwise), and $w_H(\vec{\sigma}) \equiv \sum_{i=1}^{N} \sigma_i$ is the Hamming weight of the population response $\vec{\sigma} \in \{0,1\}^N$.

We used the same learning procedure presented in Ref. [41] to compute the parameters of the Hamiltonian for the K-Pairwise Maximum Entropy model given measured constraints. The proof of convergence for the core of this L1-regularized maximum entropy algorithm is given in Ref. [59]. The code used to fit the model was written in C++ and Matlab (combined using MEX files).

Fitting the Tree Hidden Markov model

The Tree HMM models the sequence of observed population responses as (see Ref. [35]):

$$P\Big(\vec{\sigma}(1), \cdots, \vec{\sigma}(T)\Big) = \sum_{\vec{\alpha}\in[m]^T} \left(\prod_{t=1}^{T} Q_{\alpha_t}\Big(\vec{\sigma}(t)\Big)\Gamma_{\alpha_{t-1},\alpha_t}\right) \tag{3}$$

where T denotes the total number of time bins, $\vec{\alpha} \equiv (\alpha_1, \cdots, \alpha_t, \cdots, \alpha_T)$ denotes the sequence of latent mode identities, and $Q_\alpha(\cdot)$ denotes the emission distribution for mode $\alpha \in [m] \equiv \{1, \cdots, m\}$. In practice, we assume that the transition matrix Γ, which has entries $\Gamma_{\alpha_{t-1},\alpha_t} \equiv P(\alpha_t|\alpha_{t-1})$, is stationary (i.e. time-independent).

Each emission distribution Q_α in the model is a tree graphical model (see Ref. [60]). Let $G_\alpha \equiv \Big([N], \mathcal{E}_\alpha\Big)$ denote the underlying graph corresponding with Q_α, and let r_i denote the number of possible responses of neuron i. In practice, we allow the underlying graph to be a forest, so let p denote the number of connected components. Then the number of free parameters contributed by each Q_α is:

$$\begin{aligned}
\#\text{ Parameters of } Q_\alpha &= \sum_{(i,j)\in\mathcal{E}_\alpha} r_i r_j - \sum_{i\in[N]} \Big(\deg(i) - 1\Big) r_i - p \\
&= \sum_{(i,j)\in\mathcal{E}_\alpha} r_i r_j - \sum_{i\in[N]} r_i \cdot \deg(i) \\
&\quad + \sum_{i\in[N]} r_i - \underbrace{\left(N - \sum_{(i,j)\in\mathcal{E}_\alpha} 1\right)}_{p} \\
&= \sum_{(i,j)\in\mathcal{E}_\alpha} (r_i - 1)(r_j - 1) + \sum_{i\in[N]} r_i - N
\end{aligned} \tag{4}$$

In practice, we assume that all neuron responses are binary. Thus, the total number of free parameters for the entire model is:

$$\begin{aligned}
\text{Total \# Free Parameters} &= mN + m^2 + \sum_{\alpha=1}^{m} |\mathcal{E}_\alpha| \\
&\leq m^2 + 2mN - m
\end{aligned} \tag{5}$$

where $|\cdot|$ denotes cardinality. If the stationary distribution of the Markov chain is used, then the total number of free parameters reduces to $m(N+1) + \sum_{\alpha=1}^{m} |\mathcal{E}_\alpha| = \mathcal{O}(mN)$.

For a fixed number of latent modes m, the model was fit to data using the same learning procedure presented in Ref. [35]. In brief, the model parameters were inferred by maximum likelihood, using the Baum-Welch algorithm with an M-step modified to accommodate the tree graphical model form of the emission distributions

Q_α. Full details of the algorithm are described in Ref. [35]. The code used to fit the model was written in C++.

To select the number of latent modes m, also called the "latent dimensionality" [61], we carried out a 2-fold cross-validation procedure in which we randomly chose half of the time bins in the data to assign to the training set (with the other non-overlapping half assigned to the test set). For the natural movie datasets, m was then chosen to be the value which maximized the cross-validated log likelihood (CV-LL), averaged over the 2 folds. Note that to mitigate overfitting, we also incorporated a regularization parameter $\eta \in [0, 1]$ in this procedure (in practice, $\eta = 0.002$ was used throughout), as in [35]. For the white noise checkerboard dataset in the high repeat regime, in practice the CV-LL curve often began to saturate at a lower latent dimensionality than the peak. For this case, as in Lakshmanan *et al.* (2015) [61], we report the latent dimensionality at which each CV-LL curve reached 90% of its total height, where the height of each CV-LL curve is the difference between its maximum and minimum values over the range of dimensionalities tested. The motivation behind this choice is that it provides a more parsimonious representation of the data.

In the *Results* section, we also report the *normalized* cross-validated log-likelihood, $\widetilde{\mathcal{L}}$, which we define for a given latent dimensionality m' as:

$$\widetilde{\mathcal{L}}\left(m', \theta\right) = \frac{\left\langle \mathcal{L}(m', \theta) \right\rangle_n - \min_m \left[\left\langle \mathcal{L}(m, \theta) \right\rangle_n \right]}{\max_m \left[\left\langle \mathcal{L}(m, \theta) \right\rangle \right] - \min_m \left[\left\langle \mathcal{L}(m, \theta) \right\rangle \right]} \quad (6)$$

where $\mathcal{L}(m, \theta)$ denotes the log-likelihood of the test dataset $D \in \{0, 1\}^{N \times T/2}$ for the parameter values θ obtained via fitting the model to the training set, $\langle \cdot \rangle_n$ denotes the average over all n cross-validation folds, and the minimum and maximum are taken over all latent dimensionality values m considered.

Note that when exploring the probability landscape modeled by the Tree HMM (see the following section), we used the stationary distribution $\vec{\psi}$ of the fit Markov chain (assumed to be stationary), which is defined as the left eigenvector of the transition matrix Γ with unity eigenvalue (i.e. $\vec{\psi}'\Gamma = \vec{\psi}'$) and satisfies $\vec{\psi}' \cdot \vec{1} = 1$ (where $'$ denotes vector transpose). This allows us to write the modeled joint probability, $P_{\text{stationary}}$, for a single population response $\vec{\sigma} \in \{0, 1\}^N$ as:

$$P_{\text{stationary}}\left(\vec{\sigma}\right) = \sum_{\alpha=1}^{m} \psi_\alpha Q_\alpha(\vec{\sigma}) \quad (7)$$

since for a stationary Markov chain, as time $T \to \infty$, the overall proportion of time spent in mode α converges to ψ_α.

Parametric repeat analysis

For the parametric repeat analysis using Movie #3, which was an interleaved binary white-noise checkerboard movie stimulus, we normalized the length of each generated subset movie to seventy 30-s movie clips, corresponding with a total duration of 2100 s each. We quantified the "repeat ratio" as $\frac{n_r}{70}$, where n_r denotes the number of included repeats of the fixed target clip. In practice, n_r varied between 0 and 68 repeats. For each repeat ratio, the specific choice of which n_r of the 68 original repeat target clips to include, and the choice of which $(70 - n_r)$ of the 70 original unique clips to include in the corresponding subset movie, was made in accordance with one of five random sequence permutations.

For the parametric repeat analysis using Movie #4, which was an interleaved natural movie stimulus, we normalized the length of each generated subset movie to seventy-two 60-s movie clips, corresponding with a total duration of 4320 s each. In practice, n_r varied between 0 and 72 repeats. For each repeat ratio, the choice of which n_r of the 72 original repeat target clips to include, and the choice of which $(72 - n_r)$ of the 74 original unique clips to include in the corresponding subset movie, was made in accordance with one of two random sequence permutations.

Scaled count distribution

To coarsely characterize the effect of spike count on the empirical joint response probability landscape, we defined the "scaled count distribution", denoted $\widetilde{P}(K)$, as:

$$\widetilde{P}\left(K\right) \equiv \frac{P_{\text{empirical}}(K)}{\binom{N}{K}} \quad (8)$$

where

$$P_{\text{empirical}}(K) \equiv \frac{\sum_{t=1}^{T} \delta\Big(K, w_H\big(\vec{\sigma}(t)\big)\Big)}{T} \quad (9)$$

where $w_H\big(\vec{\sigma}(t)\big)$ denotes the spike count of the population response observed in time bin t, T denotes the total number of time bins in the data, and $\delta(\cdot, \cdot)$ denotes the Kronecker delta function.

Exploring the modeled probability landscape

Finding local maxima

We sought to find the (true) local maxima of the modeled probability landscape for our datasets, as defined in Ref. [41]. Note that this definition is equivalent to what

would be termed a "($\delta = 1, \rho =$ Hamming distance)-mode" in applied math [56]:

Definition 4 *A point is a* (δ, ρ)-mode *if and only if its probability is higher than all points within distance* δ *under a distance metric* ρ.

To our knowledge, with the exception of single tree graphical models, it is currently computationally intractable to find all of the local maxima of a discrete high-dimensional probability distribution. Since in their general forms neither the maximum entropy models nor the Tree Hidden Markov model are members of this tractable category, we used the same numerical procedure as in [41], which is implemented as follows: We start with a population response $\vec{\sigma}$ that appears in the data, and attempt to 'flip' the responses of neurons $i = p_1, p_2, \cdots, p_N$, where $(p_1, p_2, \cdots, p_N)$ is a randomly chosen permutation of the N neuron indices. (A different random permutation is independently generated for each starting response $\vec{\sigma}$. That is, we attempt to set $\sigma_{p_k} = \sigma_{p_k} \oplus 1$, where $\oplus$ denotes mod 2 addition. A flip is retained if $-\log P\big(\vec{\sigma}^{(p_k)}\big) < -\log P\big(\vec{\sigma}\big)$, where $\vec{\sigma}^{(p_k)}$ denotes the new response configuration obtained by flipping neuron p_k's response. When none of the N neurons can be flipped, the resulting pattern is recorded as a local maximum. We refer to this numerical procedure throughout as "single spin flip ascent." Implementation of the single spin flip ascent procedure was done in Matlab.

Finding soft local maxima

To find soft local maxima, we used a numerical procedure which we refer to throughout as "neuron pair relaxation." Conceptually, the algorithm is implemented as follows: We start with a population response pattern $\vec{\sigma}(t) \in \{0,1\}^N$ that appears in the data, and attempt to 'flip' a randomly-chosen pair (i, j) of neurons (where $i, j \in [N]$) which have opposite responses within time bin t. That is, $\sigma_i(t) = \sigma_j(t) \oplus 1$, where $\oplus$ denotes mod 2 addition. Flipping is retained if the resulting pattern has a higher probability than the original pattern. Note that the order in which these pairs are checked is chosen according to a random permutation, which is independently generated for each iteration of the algorithm. This procedure is iterated until none of the $K \cdot (N-K)$ neuron pairs with opposite instantaneous responses can be flipped to increase the probability; the terminating pattern is by definition a soft local maximum. Implementation of the neuron pair relaxation procedure was done in Matlab.

To investigate how robust the mapping from the dataset of all observed responses with K spikes to the set of unique K-soft local maxima was, we performed 100 independent iterations of our mapping procedure for the Movie #2 dataset. For each iteration, a distinct random sequence of permutations was used for the choice of neuron pairs when implementing the neuron pair relaxation algorithm. For each value of K, we then computed the mean pairwise overlap ratio between the identified set of unique K-soft local maxima for each iteration pair (l, m) (there was a total of $\binom{100}{2} = 4950$ iteration pairs), defined as:

$$\text{Mean Pairwise Overlap Ratio} \equiv \frac{\frac{1}{\binom{100}{2}} \sum_{l<m} \left|\mathcal{U}_l \cap \mathcal{U}_m\right|}{\max_{l=1}^{100} \left(\left|\mathcal{U}_l\right|\right)} \tag{10}$$

where l and m index one of the 100 mapping iterations performed, $\mathcal{U}_l$ denotes the set of unique K-soft local maxima obtained for the l-th mapping iteration, and $|\cdot|$ denotes set cardinality. Note that the the mean pairwise overlap ratio values are reported as percentages in Fig 6A.

Choice of probability model

We considered two models of the joint probability mass function (p.m.f.) of the RGC population responses: the K-Pairwise Maximum Entropy model and the Tree Hidden Markov model (recently developed, respectively, in Refs. [41] and [35]). For our purposes, it is important that the model we choose accurately captures the soft local maxima present (or not) in the empirical probability landscape. Due to limited sampling, it is intractable to determine the full empirical probability landscape and hence to make a complete comparison. However, population response states with low spike count are far better sampled than response states with high spike count. In particular, we have sufficient sampling of response states with $K = 2$ spikes. Thus, we computed the empirical $(K = 2)$-soft local maxima.

This was done by performing n iterations (in practice, we performed $n = 15$ iterations) of the following cross-validation procedure: For each iteration, we

1. Randomly split the data (i.e. the set of all responses with 2 spikes) into two training sets.

2. For each of the two training sets, we estimated the empirical p.m.f. of an observed response state $\vec{\sigma}$ as:

$$\widehat{P}_{\text{empirical}}\left(\vec{\sigma}\right) = \frac{\sum_{t=1}^{T} \delta\Big(\vec{\sigma}, \vec{\sigma}(t)\Big)}{T/2} \tag{11}$$

   where $\vec{\sigma}(t)$ denotes the population response observed in time bin t, $T/2$ was the total number of time bins in each respective training set, and $\delta(\cdot,\cdot)$ denotes the Kronecker delta function.

3. For each of the two training sets, we then performed neuron pair relaxation on all unique observed re-

sponses in the given training half, using the respective estimated empirical p.m.f. to find ($K = 2$)-soft local maxima.

To quantify the results of the above, we then computed the union of all unique empirical ($K = 2$)-soft local maxima found over the $2n$ iteration halves. Over 15 iterations (i.e. 30 different splits of the data), we found a total of 68 unique ($K = 2$)-soft local maxima. To check how reliable these results were, we then computed the proportion of occurrence of each of the 68 unique ($K = 2$)-soft local maxima, defined as follows:

$$\text{Proportion of Occurrence of } \vec{\gamma} := \frac{\sum_{h=1}^{n} \mathbb{1}_{\mathcal{S}_h}(\vec{\gamma})}{2n} \quad (12)$$

where $\vec{\gamma}$ denotes the soft local maximum of interest, $\mathcal{S}_h$ denotes the set of unique ($K = 2$)-soft local maxima found after performing anti-correlated neuron pair relaxation on the data in half h (where $1 \leq h \leq 2n$), and $\mathbb{1}$ denotes the indicator function (i.e. $\mathbb{1}_{\mathcal{S}_h}(\vec{\gamma}) = 1$ if $\vec{\gamma} \in \mathcal{S}_h$, 0 otherwise). The resulting proportions of occurrence are shown in Fig 12 (black trace).

It is possible that differences obtained between the two splits of the data in the exact identities of the empirical soft local maxima could arise from noise, attributable to limited sampling upon halving the data. To account for this possibility, we also performed a "relaxed" version of the above empirical analysis.

In the relaxed version, for each of the 68 unique empirical ($K = 2$)-soft local maxima found combined over all iterations, we performed the following: For each of the n iterations, we checked if the given ($K = 2$)-soft local maximum - which we will call $\vec{\gamma}_1$ - was found when using one data half but not the other. If so, e.g. if it was found using the first half but not the second, then we performed neuron pair relaxation on $\vec{\gamma}_1$, importantly using the p.m.f. estimated from the data in the second half. If this resulted in a ($K = 2$)-soft local maximum that had been found originally in the second half - we will call this state $\vec{\gamma}_2$ - then we say that $\vec{\gamma}_1$ and $\vec{\gamma}_2$ are equivalent. (I.e. in math terms, we identify $\vec{\gamma}_1$ and $\vec{\gamma}_2$ as members of the same equivalence class). In this case, we updated the proportion of occurrence of $\vec{\gamma}_1$ by adding a term of $\frac{1}{2n}$. Otherwise, the proportion of occurrence of $\vec{\gamma}_1$ was unmodified. The updated proportions of occurrence for the relaxed empirical analysis are shown in Fig 12 (light blue trace).

Analytical test cases

To verify that the procedures and code used in the present paper were correct and bug-free, we developed multiple test cases in which the number and identity of local maxima and/or soft local maxima is analytically known. Details are provided in the *Supporting Information*. All test cases were successfully passed (S2 and S3 Figs).

Visualizing ridges

To visualize specific examples of soft local maxima organization across low to high spike count (K) levels, we used a breadth-first search (BFS)-like algorithm. The input to this algorithm is a given "root" soft local maximum, and the output is a rooted digraph which is specific to the input. Nodes of this output rooted digraph represent soft local maxima, and a directed edge $(\vec{\gamma}, \vec{\gamma})$ exists if and only if soft local maximum $\vec{\gamma}$ is u-reachable from $\vec{\sigma}$, as defined in Definition 2. Given a fixed spike count level K (in practice we used $K = 4$), we chose each input "root" to be one of the K-soft local maximum found by performing the previously-described neuron pair relaxation procedure on the probability landscape obtained by fitting the Tree HMM to data.

Since the neuron pair relaxation algorithm involves a random permutation of the sequence of neuron pairs to flip, i.e. is not deterministic, for each neuron i in the silent set of the current K-soft local maximum, $\mathcal{S}(\vec{\gamma}_K)$, we performed 10 iterations of the neuron pair relaxation procedure on $\vec{\gamma}_K^{(i)}$, each time using a different, independent choice of the random permutation. We then set the weight of each directed edge $\left(\vec{\sigma}_K, \vec{\gamma}_{K+1}\right)$, which we denote by $w(\vec{\gamma}_K, \vec{\gamma}_{K+1})$, as the proportion of the 10 relaxation iterations which resulted in $\vec{\gamma}_{K+1}$. To mitigate effects due to noise, for each unique ($K + 1$)-soft local maximum $\vec{\gamma}_{(K+1)}$ reached after performing the above procedure on a given $\vec{\gamma}_K$, we further only included the directed edge $\left(\vec{\sigma}_K, \vec{\gamma}_{K+1}\right)$ if $w(\vec{\gamma}_K, \vec{\gamma}_{K+1})$ exceeded a reliability threshold $\Theta = 0.3$.

To aid with visualizing distinct ridges, we also assigned an x- and y-coordinate to each node $\vec{\gamma}_K$ (i.e. soft local maximum) in the output rooted digraphs as follows:

1. The y-coordinate, denoted $y(\vec{\gamma}_K)$, was set to K (i.e. the spike-count of $\vec{\gamma}_K$).

2. Our choice of formulation for the x-coordinate was based on two motivating criteria: First, we wanted to be able to visualize potential "jumps" in ridge organization, due to the current identified ridge either terminating at a lower spike count level than the one we chose arbitrarily to explore out to, or/and branching into multiple distinct ridges. Examples of this scenario can be seen in Figs 8C and 9B. Second, we wanted to visualize only "genuine" jumps or branching, with respect to the notion of neuron communities. Specifically, after inspection of preliminary digraph visualizations, we noticed that 'connected' (i.e. u-reachable) soft local

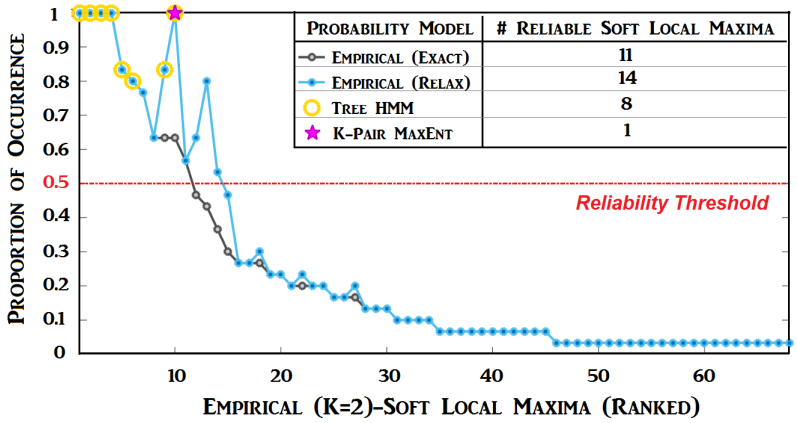

FIG. S12: **Comparison of different models with the empirical (K=2)-soft local maxima results. (A)** Shown in (blue) grey is the (relaxed) proportion of occurrence of each of the 68 ($K = 2$)-soft local maxima found using the halving procedure with the empirical probabilities (ranked). The number of ($K = 2$)-soft local maxima captured by each model with a proportion of occurrence greater than the chosen reliability threshold of $\Theta = 0.5$ is shown in the table inset. The Tree HMM results (denoted by yellow annuli) matched the empirical results the best, capturing 8 of the reliable empirical ($K = 2$)-soft local maxima. In contrast, the K-Pairwise Maximum Entropy model (purple star) captured only one ($K = 2$)-soft local maximum.

maxima seemed to each have an activity set of active neurons that was a subset of a larger unique group (which we term the "neuron community"). We thus post-hoc formulated our definition of the x-coordinate for our digraph visualizations to aid us in examining how prevalent this property was. Since community structure is a clustering property in the graph of neurons, we needed to define an overlap measure between the activity sets of different soft local maxima. The natural distance metric based on overlap is:

$$d_O\left(\vec{\gamma}_K, \vec{\gamma}_G\right) := \left|\mathcal{A}\left(\vec{\gamma}_K\right) \cap \mathcal{A}\left(\vec{\gamma}_G\right)\right| - K \tag{13}$$

where $G > K$ denotes a higher spike-count level, $\mathcal{A}(\vec{\gamma}) := \{i|\gamma_i = 1\}$ denotes the "activity set" (i.e. set of active neurons) of soft local maximum $\vec{\gamma}$, and $|\cdot|$ denotes cardinality. Note that if there is complete overlap of the activity sets, i.e. $\mathcal{A}(\vec{\gamma}_K) \subset \mathcal{A}(\vec{\gamma}_G)$, then $\left|\mathcal{A}\left(\vec{\gamma}_K\right) \cap \mathcal{A}\left(\vec{\gamma}_G\right)\right| = K$, and thus $d_O\left(\vec{\gamma}_K, \vec{\gamma}_G\right) = 0$.

Conceptually, for a given $\vec{\gamma}_K$, we first computed the succession of all pairwise u-reachable soft local maxima out to the chosen maximum spike count level, and allowed the algorithm to potentially consider all soft local maxima $\vec{\gamma}_G$ at a higher spike count level (i.e. $G > K$). We then chose to assign to $\vec{\gamma}_K$ the same x-coordinate as the soft local maximum which was closest in terms of overlap, $\vec{\gamma}^*_G$, plus an offset equal to the overlap distance between $\vec{\gamma}_K$ and $\vec{\gamma}^*_G$. That is, formally, we assigned:

$$x\left(\vec{\gamma}_K\right) := x\left(\vec{\gamma}^*_G\right) + d_O\left(\vec{\gamma}_K, \vec{\gamma}^*_G\right) \tag{14}$$

where

$$\vec{\gamma}^*_G := \underset{\vec{\gamma}_G\,:\,G>K}{\arg\min}\, d_O\left(\vec{\gamma}_K, \vec{\gamma}_G\right) \tag{15}$$

In the case of a tie for Eq. 14 (for example, if the considered $\vec{\gamma}_4$ had complete overlap of its activity set with both $\vec{\gamma}_5$ and $\vec{\gamma}_6$), we chose the soft local maximum at the nearest spike count level to be $\vec{\gamma}^*_G$ (i.e. in the above example, we would choose $\vec{\gamma}^*_G = \vec{\gamma}_5$, and would assign $\vec{\gamma}_4$ the same x-coordinate as $\vec{\gamma}_5$).

Fig 8A provides an example of why this formulation considering all higher spike count levels is advantageous for our purposes of assigning the x-coordinate based on only pairwise overlap: If the latter were used, then $\vec{\gamma}_9$ in Fig 8A would be assigned an x-coordinate offset from $\vec{\gamma}_8$ by 1 unit, resulting in a new branch. However, this branch would be artificial with respect to the notion of a neuron community, since $\mathcal{A}(\vec{\gamma}_8) \subset \mathcal{A}(\vec{\gamma}_{10})$ and $\mathcal{A}(\vec{\gamma}_9) \subset \mathcal{A}(\vec{\gamma}_{10})$. Thus, $\vec{\gamma}_8$ should be assigned the same x-coordinate at $\vec{\gamma}_{10}$.

In addition, we assigned a color to each node in each output rooted digraph visualization, which uniquely corresponds with the MAP (maximum *a posteriori*) estimate α of the latent Tree HMM mode for the associated soft local maximum $\vec{\gamma}$, given by:

$$\begin{aligned}\hat{\alpha}_{MAP}(\vec{\gamma}) &\equiv \arg\max_{\alpha}\left[\frac{P(\vec{\gamma}|\alpha)}{\int_\beta P(\vec{\gamma}|\beta)P(\beta)d\beta}\right] \\ &= \arg\max_{\alpha}\left[\psi_\alpha Q_\alpha(\vec{\gamma})\right]\end{aligned} \tag{16}$$

where ψ denotes the stationary distribution of the Markov chain, and $Q_\alpha(\cdot)$ is the emission distribution for

mode α (see *Methods: Fitting the Tree Hidden Markov Model*). For the dataset of $N = 152$ ganglion cells responding to the non-repeated natural movie (i.e. Movie #2), the optimal latent dimensionality was 19 Tree HMM modes; there were thus 19 unique possible colors. Visualization of the rooted digraphs computed via the BFS algorithm (examples of which are shown in Figs 8 and 9) was automated using the *igraph* package in R.

We also investigated the organization of soft local maxima across high to low spike count (K) levels; example results are shown in S7 Fig. For this analysis, a rooted digraph was constructed, but where each root node was taken to be a soft local maximum found (by performing neuron pair relaxation on the data) at a high K level (e.g. $K = 11$), and where a directed edge $(\vec{\sigma}, \vec{\gamma})$ was added if and only if soft local maximum $\vec{\gamma}$ is d-reachable from $\vec{\sigma}$ (see Definition 3). Note that we report the u-reachable version (i.e. starting from a low spike count root node and progressing out to higher spike count levels) in the main text, as we have more reliable sampling of the data at low spike count levels.

Visualizing neuron communities

Fig 10A was made using the *igraph* package in R.

Shuffled Controls

Heterogeneous Firing Rates

To generate this control, for each neuron i in the population we performed a "complete shuffle" on neuron i's discretized spike train (binned into T total 20 ms time bins) occurring in the data, by implementing a random permutation of the T time bins. Note that an independent random permutation was implemented for each neuron. This complete shuffling procedure eliminates all signal and noise correlations among neurons in the population, but retains the firing rate of each neuron over the course of the experimental recording. Thus, the control 'dataset' of joint responses generated via this procedure corresponds with a population of neurons which fire independently, but which have heterogeneous firing rates that match the original data.

Homogeneous Firing Rates

To generate this control, we first computed the average firing rate from the data, $\langle r_i \rangle_{i=1}^{N}$, over the entire population of N ganglion cells, where $r_i \equiv \left(\sum_{t=1}^{T} \sigma_i(t) \right) / T$. We then simulated an independent population of matched size N in which each neuron i was assigned to have its firing rate $r_i = \rho \equiv \langle r_i \rangle_{i=1}^{N}$. Thus, the control 'dataset' of joint responses generated via this procedure corresponds with a population of neurons which fire independently, and moreover have homogeneous firing rates.

[1] Perkel D, Bullock T. Neural coding. Neurosci Res Program Bull. 1968;6:221343.

[2] Rieke F, Warland D, de Ruyter van Steveninck R, Bialek W. Spikes: exploring the neural code. MIT Press; 1999.

[3] Dayan P, Abbott L. Theoretical neuroscience. MIT Press; 2001.

[4] Kalaska JF, Crammond DJ. Cerebral cortical mechanisms of reaching movements. Science. 1992;255:1517-23.

[5] Ahissar M, Ahissar E, Bergman H, Vaadia E. Encoding of sound-source location and movement: activity of single neurons and interactions between adjacent neurons in the monkey auditory cortex. J Neurophysiol. 1992;67:203-215.

[6] Georgopoulos AP. Current isses in directional motor control. Trends Neurosci. 1995;18:506-510.

[7] Nicolelis MAL, Baccala LA, Lin RC, Chapin JK. Sensorimotor encoding by synchronous neural ensemble activity at multiple levels of the somatosensory system. Science. 1995;268:13531358.

[8] Sakurai Y. Population coding by cell assemblies - what it really is in the brain. Neurosci Res. 1996;26:1-16.

[9] Ghazanfar AA, Nicolelis MAL. Non-linear processing of tactile information by thalamocortical ensembles. J Neurophysiol. 1997;78:506-510.

[10] Nicolelis MAL, Fanselow EE, Ghazanfar AA. Hebb's dream: The resurgence of cell assemblies. Neuron. 1997;19:219-21.

[11] Stopfer M, Vivek J, Gilles L. Intensity versus identity coding in an olfactory system. Neuron. 2003;39(6):991-1004.

[12] Puchalla JL, Schneidman E, Harris RA, Berry MJ. Redundancy in the population code of the retina. Neuron. 2005; 46(3):493-504.

[13] Mazor O, Laurent G. Transient dynamics versus fixed points in odor representations by Locust antennal lobe projection neurons. Neuron. 2005;48:661673.

[14] Fujisawa S, Amarasingham A, Harrison MT, Buzsaki G. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. Nat Neurosci. 2008; 11:823833.

[15] Truccolo W, Hochberg LR, Donoghue JP. Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. Nat Neurosci. 2010;13:105111.

[16] Harvey CD, Coen P, Tank DW. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature. 2012;484:6268.

[17] Ganmor E, Segev R, Schneidman E. A thesaurus for a neural population code. Elife. 2015;4:e06134.

[18] Greenberg DS, Houweling AR, Kerr JND. Population imaging of ongoing neuronal activity in the visual cortex of awake rats. Nat Neurosci. 2008;11:749-751.

[19] Shlens J, Field GD, Gauthier JL, Greschner M, Sher A, Litke AM, Chichilnisky EJ. The structure of large-scale synchronized firing in primate retina. J Neurosci. 2009

Apr 15;29(15):5022-31.
[20] Marre O, Amodei D, Deshmukh N, Sadeghi K, Soo F, Holy TE, *et al.* Mapping a complete neural population in the retina. J Neurosci. 2012 Oct 24;32(43):14859-38.
[21] Ahrens MB, Orger MB, Robson DN, Li JM, Keller PJ. Whole-brain functional imaging at cellular resolution using light-sheet microscopy. Nat Methods. 2013 March 12;10:413-20.
[22] Schneidman E, Berry MJ II, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. Nature. 2006;440:10071012.
[23] Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP. Spatio-temporal correlations and visual signalling in a complete neuronal population. Nature. 2008;454:995999.
[24] Osborne LC, Palmer SE, Lisberger SG, Bialek W. The Neural Basis for Combinatorial Coding in a Cortical Population Response. J Neurosci. 2008 Dec 10;28(50):13522-31.
[25] Marre O, Boustani EI, Fregnac Y, Destexhe A. Prediction of Spatiotemporal Patterns of Neural Activity from Pairwise Correlations. Phys Rev Lett. 2009 April 2;102(13):138101.
[26] Ohiorhenuan IE, Mechler F, Purpura KP, Schmid AM, Hu Q, Victor JD. Sparse coding and high-order correlations in fine-scale cortical networks. Nature. 2010 Jul 29;466(7306):617-21.
[27] Ganmor E, Segev R, Schneidman E. Sparse low-order interaction network underlies a highly correlated and learnable neural population code. Proc Natl Acad Sci USA. 2011;108:96799684.
[28] Schwartz G, Macke J, Amodei D, Tang H, Berry MJ, 2nd. Low error discrimination using a correlated population code. J Neurophysiol. 2012;108:10691088.
[29] Zohary E, Shadlen MN, Newsome WT. Correlated neuronal discharge rate and its implications for psychophysical performance. Nature. 1994;370:140143.
[30] Abbott LF, Dayan P. The effect of correlated variability on the accuracy of a population code. Neural Comput. 1999;11:91101.
[31] Sompolinsky H, Yoon H, Kang K, Shamir M. Population coding in neuronal systems with correlated noise. Phys Rev E. 2001 Oct 17;64:051904.
[32] da Silveira RA, Berry MJ, 2nd. High-fidelity coding with correlated neurons. PLoS Comput Biol. 2014 Nov 20;10(11):e1003970.
[33] Vidne M, Ahmadian Y, Shlens J, Pillow JW, Kulkarni J, Litke AM, Chichilnisky EJ, Simoncelli E, Paninski L. Modeling the impact of common noise inputs on the network activity of retinal ganglion cells. J Comput Neurosci. 2012;33(1):97-121.
[34] Granot-Atedgi E, Tkacik G, Segev R, Schneidman E. Stimulus-dependent maximum entropy models of neural population codes. PLoS Comput Biol. 2013;9(3):e1002922.
[35] Prentice JS, Marre O, Ioffe ML, Tkacik G, Berry MJ, 2nd. Error-robust modes of the retinal population code. PLoS Comput Biol. 2016. In press.
[36] Shlens J, Field GD, Gauthier JL, Grivich MI, Petrusca D, Sher A, Litke AM, Chichilnisky EJ. The structure of multi-neuron firing patterns in primate retina. J Neurosci. 2006 Aug 9;26(32):8254-66.
[37] Vasquez JC, O, Palacios AG, Berry MJ, 2nd, Cessac B. Gibbs distribution analysis of temporal correlations structure in retina ganglion cells. arXiv;2011:1112.2464v1.
[38] Ganmor E, Segev R, Schneidman E. The Architecture of Functional Interaction Networks in the Retina. J Neurosci. 2011 Feb 23;31(8):3044-3054.
[39] Nasser H, Marre O, Cessac B. Spatio-temporal spike train analysis for large scale networks using the maximum entropy principle and Monte Carlo method. J Stat Mech. 2013 March 12;03P03006.
[40] Watanabe T, Hirose S, Wada H, Imai Y, Machida T, Shirouzu I, Konishi S, Miyashita Y, Masuda N. A pairwise maximum entropy model accurately describes resting-state human brain networks. Nat Commun. 2013 Jan 22;1370.
[41] Tkacik G, Marre O, Amodei D, Schneidman E, Bialek W, Berry MJ, 2nd. Searching for collective behavior in a large network of sensory neurons. PLoS Comput Biol. 2014 Jan;10(1):e1003408.
[42] Mora T, Deny S, Marre O. Dynamical Criticality in the Collective Activity of a Population of Retinal Neurons. Phys Rev Lett. 2015 Feb 20;114:078105.
[43] Parisi G, Mezard M, Virasoro MA. Spin glass theory and beyond. World Scientific; 1987.
[44] Mainen ZF, Sejnowski TJ. Reliability of spike timing in neocortical neurons. Science. 1995;268:15031506.
[45] Schneidman E, Freedman B, Segev I. Ion channel stochasticity may be critical in determining the reliability and precision of spike timing. Neural Comput. 1998;10:16791703.
[46] Osborne LC, Lisberger SG, Bialek W. A sensory source for motor variation. Nature. 2005;437:412-6.
[47] Faisal AA, Selen LP, Wolpert DM. Noise in the nervous system. Nat Rev Neurosci. 2008;9:292303.
[48] Berry MJ, Warland DK, Meister M. The structure and precision of retinal spike trains. Proc Natl Acad Sci USA. 1997;94:54115416.
[49] de Ruyter van Steveninck RR, Lewen GD, Strong SP, Koberle R, Bialek W. Reproducibility and variability in neural spike trains. Science. 1997;275:18051808.
[50] Reich DS, Mechler F, Victor JD. Independent and redundant information in nearby cortical neurons. Science. 2001;294:25662568.
[51] Shannon, CE. A mathematical theory of communication. Bell Labs Tech J. 1948; 27:379423, 623656.
[52] Roth, RM. Introduction to Coding Theory. Cambridge University Press; 2006.
[53] Sreenivasan S, Fiete I. Grid cells generate an analog error-correcting code for singularly precise neural computation. Nat Neurosci. 2011;14:13301337.
[54] Curto C, Itskov V, Morrison K, Roth Z, Walker JL. Combinatorial neural codes from a mathematical coding theory perspective. Neural Comput. 2013;25:18911925.
[55] Fortunato S. Community detection in graphs. Phys Rep. 2010;485(3):75-132.
[56] Chen C, Liu H, Metaxas DN, Zhao T. Mode Estimation for High Dimensional Discrete Tree Graphical Models. Adv Neural Inf Process Syst. 2014 Dec;27:5533-51.
[57] Haralick R. Ridges and valleys in digital images. CVGIP. 1983;22:28-38.
[58] Eberly D, Gardner R, Morse B, Pizer S, Scharlach C. Ridges for image analysis. J Math Imaging Vis. 1994;4:353-373.
[59] Dudik M, Phillips SJ, Schapire RE. Performance guarantees for regularized maximum entropy density estimation.

Proceedings 17th Annual conference on learning theory. 2004.
[60] Meila M, Jaakkola T. Tractable Bayesian learning of tree belief networks. Stat Comput. 2006 Jan;16(1):7792.
[61] Lakshmanan KC, Sadtler PT, Tyler-Kabara EC, Batista AP, Yu BM. Extracting low-dimensional latent structure from time series in the presence of delays. Neural Comput. 2015;27:1825-56.
[62] Schneidman E, Puchalla JL, Segev R, Harris RA, Bialek W, Berry MJ, 2nd. Synergy from silence in a combinatorial neural code. J Neurosci. Nov 2;31(44):1573215741.
[63] Huang H, Toyoizumi T. Clustering of neural code words revealed by a first-order phase transition. Phys Rev E. 2016;93.
[64] Hebb DO. The Organization of Behavior. New York: Wiley and Sons; 1949.
[65] Wickelgren WA. Webs, cell assemblies, and chunking in neural nets. Concepts Neurosci. 1992;1:1-53.
[66] Legéndy CR. On the scheme by which the human brain stores information. Math Biosci. 1967;1:555-597.
[67] Kanerva P. Sparse Distributed Memory. Cambridge: MIT Press; 1988.
[68] Barlow, HB. Single units and sensation: a doctrine for perceptual psychology? Perception. 1972;1:371-394.
[69] Brenner N, Strong SP, Koberle R, Bialek W, de Ruyter van Stevenick RR. Synergy in a neural code. Neural Comput. 2000;12:1531-1552.
[70] Quian Quiroga R, Panzeri S. Extracting information from neuronal populations: information theory and decoding approaches. Nat Rev Neurosci. 2009 Mar;10(3):173-85.
[71] Schnitzer MJ, Meister M. Multineuronal firing patterns in the signal from eye to brain. Neuron. 2003 Feb;37:499-511.
[72] Mezard M, Montanari A. Information, Physics, and Computation. Oxford University Press, Oxford; 2009.
[73] Dong DW, Atick JJ. Statistics of natural time-varying images. Network. 1995;6:345358.
[74] Victor JD, Purpura KP. Nature and precision of temporal coding in visual cortex: A metric-space analysis. J Neurophysiol. 1996 Aug;76(2):1310-27.
[75] Victor JD, Purpura KP. Spatial phase and the temporal structure of the response to gratings in V1. J Neurophysiol. 1998 Aug;80(2):554-71.
[76] Meister M, Berry MJ, 2nd. The neural code of the retina. Neuron. 1999 Mar;22:435450.
[77] Fairhall AL, Lewen JD, Bialek W, de Ruyter van Stevenick RR. Efficiency and ambiguity in an adaptive neural code. Nature. 2001 Aug;412:787-92.
[78] Ayaz A, Saleem AB, Scholvinck ML, Carandini M. Locomotion controls spatial integration in mouse visual cortex. Curr Biol. 2013;23:890894.

SUPPLEMENTARY INFORMATION

§1. Analytical Test Cases

To verify that the procedures presented in the main text were correct, and that the corresponding code was bug-free, we developed three test cases in which the number of local maxima and/or soft local maxima was analytically known. A detailed description of each analytical test case is provided in the following.

Test Case I

The number and identity of all true local maxima in the probability landscape is analytically known for this test case, which we define as follows: Let $\mathcal{S} = \{s_1, s_2, s_3\}$ denote the stimulus space, where $s_i \sim \mathcal{U}\{1, 3\}$, i.e. denotes a random variable which is uniformly distributed. The scenario is a population of N neurons, where each neuron i has its firing rate r_i defined as

$$r_i[s_l] = \begin{cases} c & , \, i \in Z_l \subset [N] \\ 0 & , \, o.w. \end{cases} \tag{17}$$

where $l \in [3]$ (i.e. l indexes stimuli and i indexes neurons), and $r_i[s_l] = 0 \; \forall \, i \in [N] \setminus \cup_l Z_l$. Without loss of generality, we take $Z_1 \equiv \{1, 3, 5\}$, $Z_2 \equiv \{9, 11, 13\}$, $Z_3 \equiv \{17, 19, 21\}$, and take the constant $c = 0.9$. Let $\mathcal{M} \equiv \{M_1, M_2, M_3\}$, where $M_1 = (101010 \cdots 0)$, $M_2 = (00000000101010 \cdots 0)$, and $M_3 = (\underbrace{0 \cdots 0}_{16} 101010 \cdots 0)$. Note that neurons in this scenario fire independently conditioned on the stimulus, i.e.

$$P[\vec{\sigma}] = \frac{1}{3} \sum_{l=1}^{3} \left[\prod_{i \in [N]} r_i[s_l]^{\sigma_i} (1 - r_i[s_l])^{(1-\sigma_i)} \right] \tag{18}$$

Thus, combining Eqs. 17 and 18,

$$P[\vec{\sigma}] = \begin{cases} \frac{1}{3} \cdot c^3 & , \, \vec{\sigma} \in \mathcal{M} \\ \frac{1}{3} \left[c^2 (1-c) \right] & , \, \exists \; l \text{ s.t. } \{i | \sigma_i = 1\} \cap Z_l = 2, \;\; \& \;\; w_H(\vec{\sigma}) = 2 \\ \frac{1}{3} \left[c (1-c)^2 \right] & , \, \exists \; l \text{ s.t.} \{i | \sigma_i = 1\} \cap Z_l = 1, \;\; \& \;\; w_H(\vec{\sigma}) = 1 \\ (1-c)^3 & , \, \vec{\sigma} = \vec{0} \\ 0 & , \, otherwise \end{cases} \tag{19}$$

where $\vec{\sigma} \in \{0, 1\}^N$ denotes a possible population response pattern, and $w_H(\cdot)$ denotes Hamming weight.

Note that $\forall \; i \in [N] \setminus (Z_1 \cup Z_2 \cup Z_3)$, a state with $\sigma_i = 1$ will introduce a 0 in the product for all three summand terms in Eq. (2), which is why $P[\vec{\sigma}] = 0$ for all other states than those listed in Eq. (3). Thus, if a Hamming distance of one is used as the distance metric, then there are exactly three local maxima: M_1, M_2 and M_3. For illustrative purposes, an example possible energy descent trajectory is (where below we use the index set notation for a response state, i.e. the bijective mapping $\vec{\sigma} \in \{0, 1\}_2^N \leftrightarrow \{i \in [N] | \sigma_i = 1\}$):

$$\{2\} \longrightarrow \emptyset \longrightarrow \{1\} \longrightarrow \{1, 3\} \longrightarrow \{1, 3, 5\} \equiv M_1$$

Test Case I was successfully passed by the combination of our Tree HMM fitting code and single spin flip ascent code (see S2A Fig for results).

Test Case II

The number and identity of all local maxima in the probability landscape is analytically known for this test case, defined as follows: Let $\mathcal{S} = \{s_1, s_2, s_3, s_4\}$ denote the stimulus space, where $s_i \sim \mathcal{U}\{1, 4\}$. The scenario is a population

of N neurons, where each neuron i's firing rate r_i is given by:

$$\begin{aligned} r_i[s_1] &= \begin{cases} c_1 = 0.7 & , \, i \in [20] \\ 0 & , \, 21 \leq i \leq N \end{cases} \\ r_i[s_2] &= \begin{cases} 0 & , \, i \in [20] \\ c_1 = 0.7 & , \, 21 \leq i \leq N \end{cases} \\ r_i[s_4] &= \begin{cases} c_2 = 0.8 & , \, i \in \{1, N\} \\ 0 & , \, otherwise \end{cases} \end{aligned} \tag{20}$$

and $r_i[s_3] = 0 \, \forall i \in [N]$. There are exactly four analytically-known local maxima: $\vec{0}$, $\{1, 2, \cdots, 20\}$, $\{21, 22, \cdots, N\}$, and $\{1, N\}$ (where we have represented each binary response pattern by its set of active neurons). It can be shown analytically that for each of these four response states, changing any one neuron's response will result in a response state with higher energy than the original state.

Test Case II was successfully passed by the combination of our Tree HMM fitting code and single spin flip ascent code (see S2B Fig for results).

Test Case III

In this test case, all possible ($K = 4$)-soft local maxima of the probability landscape are analytically known. The scenario is a stimulus space, $\mathcal{S}$, consisting of six uniformly-presented stimuli, i.e. $\mathcal{S} \equiv \{s_l\}_{l=1}^{6}$, where $s_l \sim \mathcal{U}\{1, 6\}$. Each stimulus s_l maps to one of six population response vectors of Hamming weight $w_H(\vec{\sigma}) = 6$ or 0 in the population response space, $\{0, 1\}^N$ (where here $N = 30$ neurons). I.e. the population of 30 neurons is partitioned into six "communities" of subsets of neurons. In particular, this toy system uses the following encoding:

$$\begin{aligned} s_1 &\rightarrow c_1 \equiv \emptyset \\ s_2 &\rightarrow c_2 \equiv \{1, 2, 3, 4, 5, 6\} \\ s_3 &\rightarrow c_3 \equiv \{7, 8, 9, 10, 11, 12\} \\ s_4 &\rightarrow c_4 \equiv \{13, 14, 15, 16, 17, 18\} \\ s_5 &\rightarrow c_5 \equiv \{19, 20, 21, 22, 23, 24\} \\ s_6 &\rightarrow c_6 \equiv \{25, 26, 27, 28, 29, 30\} \end{aligned} \tag{21}$$

Noise is introduced according to a Binary Asymmetric Channel (BAC), where $\forall \, i \in [N]$,

$$r_i[s_l] \equiv P[\sigma_i = 1 | s_l] = \begin{cases} p \equiv 0.2 & , \, i \in c_l \\ q \equiv 0.001 & , \, o.w. \end{cases} \tag{22}$$

In this scenario, neuron responses are conditionally independent given the stimuli, i.e.

$$P[\vec{\sigma}] = \frac{1}{6} \sum_{l=1}^{6} \left[\prod_{i=1}^{30} r_i[s_l]^{\sigma_i} (1 - r_i[s_l])^{(1-\sigma_i)} \right] \tag{23}$$

In this case, there is only one true local maxima: the all-silent state, $\vec{0} \equiv c_1$. However, if we restrict ourselves to only looking at population responses with spike count K (e.g. $K = 4$), then any state $\vec{\sigma} \in \{0, 1\}^N$ which satisfies $\vec{\sigma} \subseteq c_l$ for some $2 \leq l \leq 6$ will be stable under anti-correlated neuron pair relaxation dynamics (i.e. will be a K-soft local maximum).

Test Case III was successfully passed by the combination of our Tree HMM fitting code and anti-correlated neuron pair relaxation code (see S3 Fig for results).

§2. SUPPLEMENTAL FIGURES

§2.a. Supplementary Figure 1.

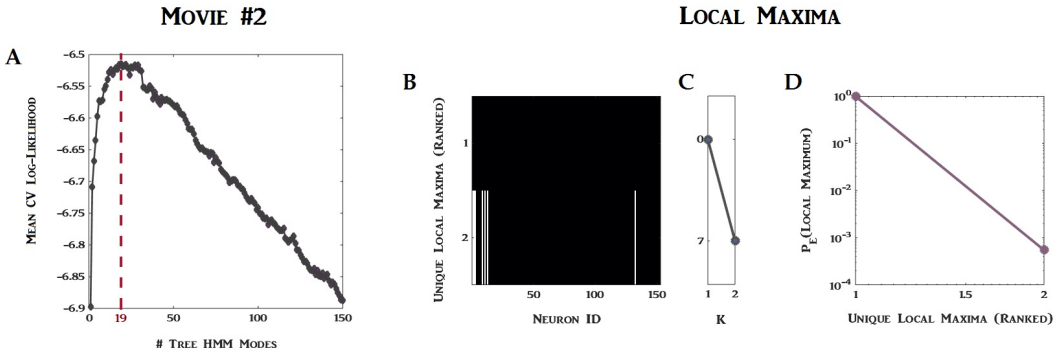

FIG. S13: **Local maxima results for the dataset of 152 RGCs responding to Movie #2. (A)** Cross-validated log-likelihood averaged over the two cross-validation folds (y-axis), as a function of the number of latent Tree HMM modes (x-axis). Red dashed line denotes the optimal latent dimensionality, which corresponds with the peak of the CV-LL curve. **(B)** Each row represents the binary representation of an identified unique soft local maximum. **(C)** The spike count of each associated soft local maximum shown in panel (B). **(D)** Log-log plot of the proportion (denoted by P_E) of the 90,001 population responses observed in the data which were mapped via the single spin flip ascent procedure to the corresponding soft local maximum, indicated on the x-axis (ranked).

§2.b. Supplementary Figure 2.

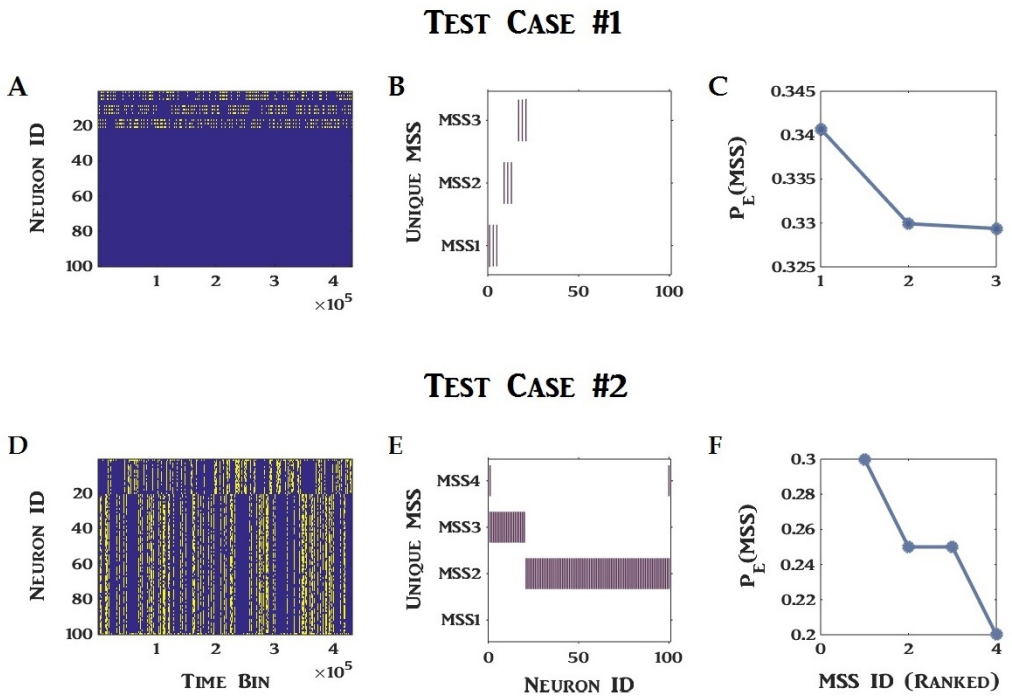

FIG. S14: **Results for analytical test cases I and II. (A-C)** Results for Test Case I. **(A)** Shown by each column is the binary representation of one of the 400,000 simulated population responses for Test Case I (see S1 Appendix for details). Blue denotes silence; yellow denotes activity. **(B)** Shown by each row is the binary representation (ticks denote spiking; white denotes silence) of each of the three unique local maxima identified via performing our single spin flip ascent procedure on the simulated dataset in (A). **(C)** Log-log plot of the proportion (denoted by P_E) of the 400,000 simulated population responses which were mapped via single spin flip ascent to the corresponding local maximum, indicated on the x-axis (ranked). **(D-F)** Results for Test Case II (see S1 Appendix for details); format is the same as is panels (A-C).

§2.c. Supplementary Figure 3.

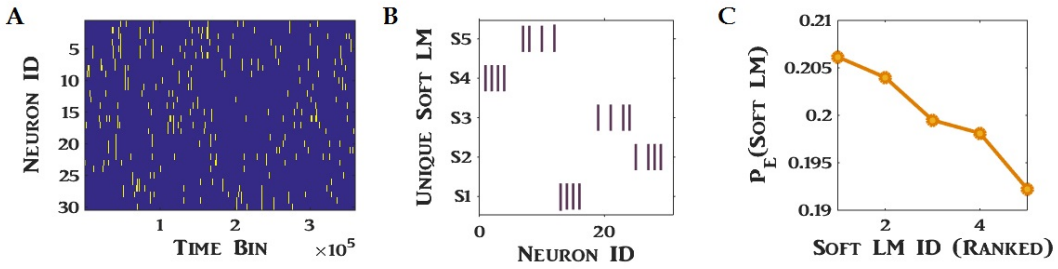

FIG. S15: **Results for analytical test case III. (A)** Shown by each column is the binary representation of one of the 400,000 simulated population responses for Test Case III (see S1 Appendix for details). Blue denotes silence; yellow denotes activity. **(B)** Shown by each row is the binary representation (ticks denote spiking; white denotes silence) of each of the five unique soft local maxima identified via performing our neuron pair relaxation procedure on the simulated dataset in (A). **(C)** Log-log plot of the proportion (denoted by P_E) of the 400,000 simulated population responses which were mapped via neuron pair relaxation to the corresponding soft local maximum, indicated on the x-axis (ranked).

§2.d. Supplementary Figure 4.

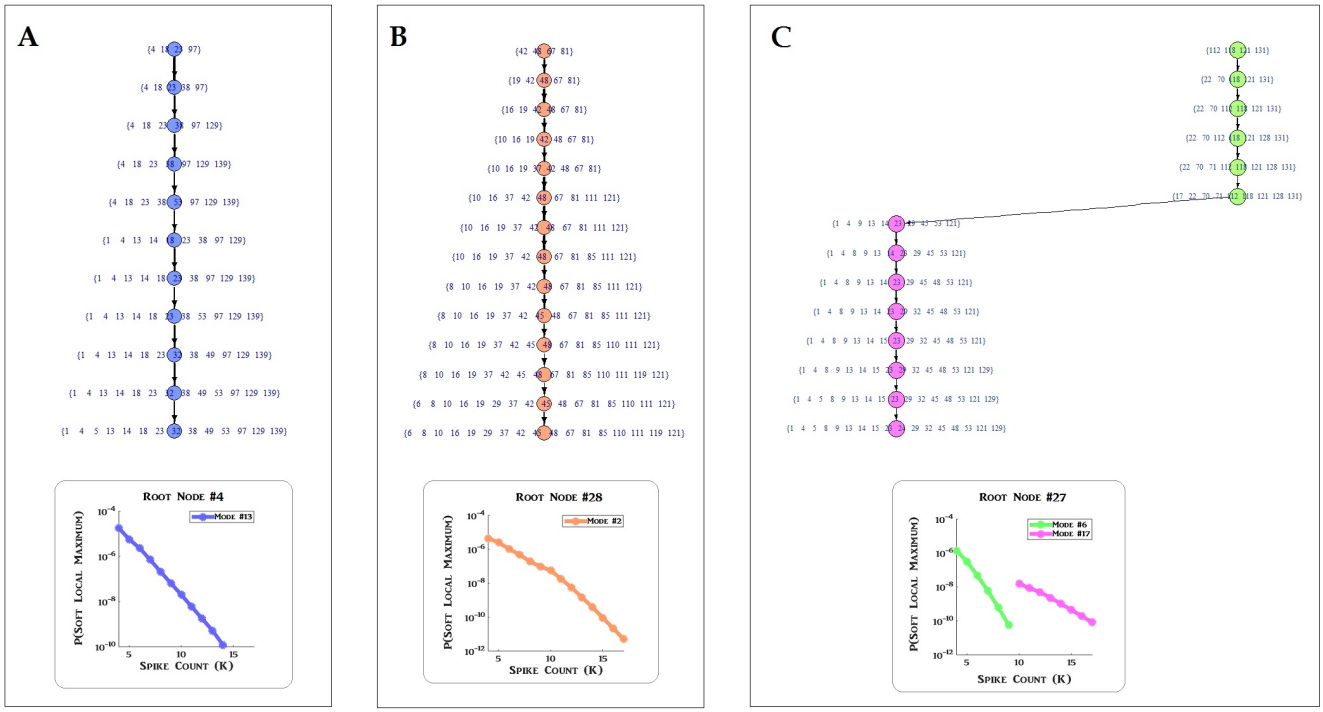

FIG. S16: **Type I ridge results, including activity sets.** Same as Fig 8, except that the full activity set $\mathcal{A}(\vec{\gamma})$ for each node $\vec{\gamma}$ in the three rooted digraph examples is written out in full.

§2.e. Supplementary Figure 5.

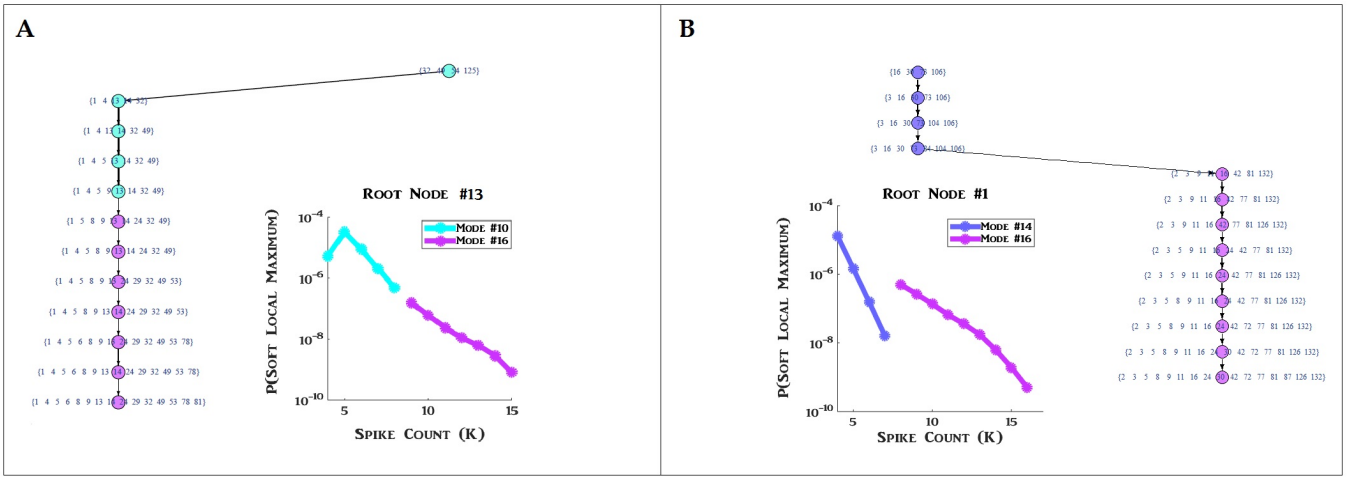

FIG. S17: **Type II ridge results, including activity sets.** Same as Fig 9, except that the full activity set $\mathcal{A}(\vec{\gamma})$ for each node $\vec{\gamma}$ in the two rooted digraph examples is written out in full.

§2.f. Supplementary Figure 6.

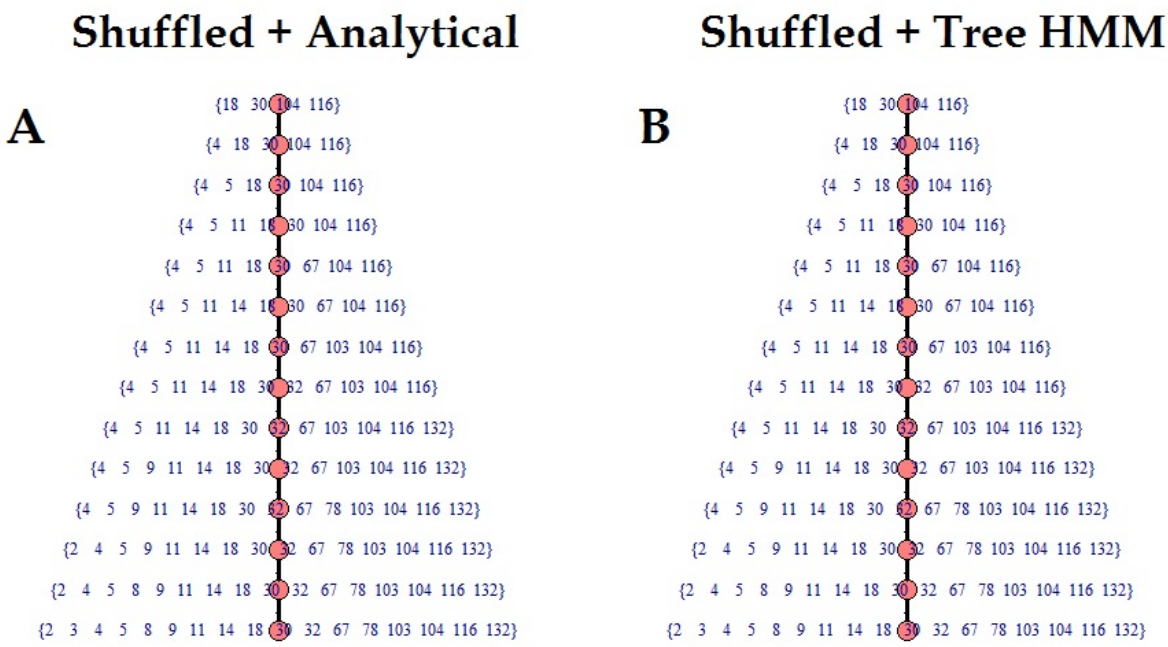

FIG. S18: **Ridge results for the heterogeneous firing rate shuffled control. (A)** Result obtained when the analytical independent model, given by $P_A(\vec{\sigma}) = \prod_i r_i^{\sigma_i}(1-r_i)^{(1-\sigma_i)}$, is used as the underlying probability model. **(B)** Result obtained when the Tree HMM is used as the underlying probability model.

§2.g. Supplementary Figure 7.

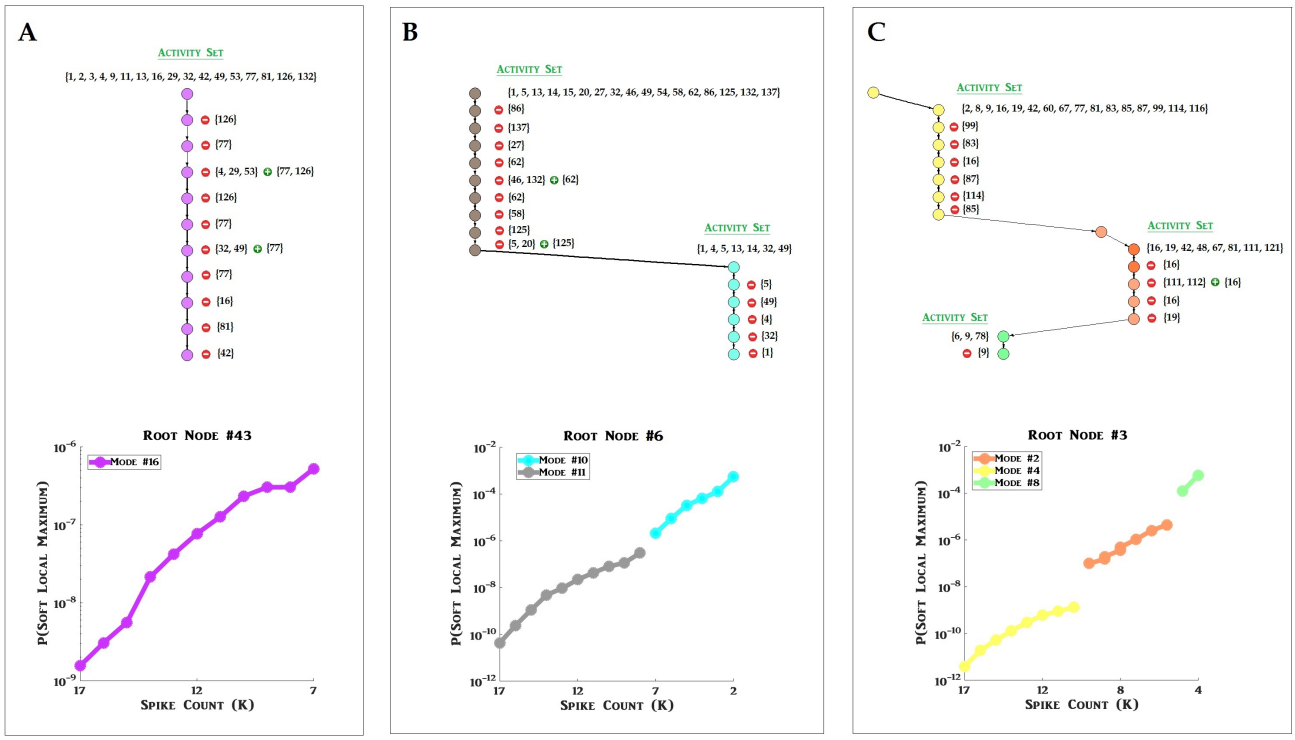

FIG. S19: **Examples of output rooted digraphs using d-reachable edges.** **(A)** Output rooted weighted digraph obtained via our BFS procedure, described in the main text, when we took the unique ($K = 17$)-soft local maximum with rank 43 as the input root node. Notation is the same as in Figs 8 and 9. The rooted weighted digraph obtained when the ($K = 17$)-soft local maximum with rank 6 and 3 was taken as the input root node is shown in **(B)** and **(C)**, respectively.